Intel Developer Update is Intel's monthly online news magazine for developers. As the official publication of developer.intel.com, it brings hardware, software, and Web developers the latest information on Intel initiatives, technologies, platforms, and products.

Cover Story
Each month, we run a cover story on the most significant industry announcement, trend, or development for the month.

Top Stories
The issue's Top Stories are the three or four articles that cover the month's hottest topics. For a list of all the articles in the issue, see the Table of Contents.

Contact the Editor
To make *Intel Developer Update* a better information resource, we invite you to share your thoughts on what we've published or what you'd like to see covered. Comments are always welcome.

Archives
Our archives contain two groups of previously published articles. One group contains all the articles that appeared in *Platform Solutions News*, the earlier version of *Intel Developer Update*. The articles date from September 1997 through August 1999. The other group is set up to contain *Intel Developer Update* articles dating from the inaugural September/October 1999 issue.

Bookmarking
We advise against bookmarking article pages. They're accessible online only during the month the issue is live. Thereafter, they're removed to our archives. Instead, we suggest that you bookmark the PDF (Adobe* Portable Document Format) file versions of the articles. You'll find buttons for the PDF files labeled "print article" in the right navigation section of each article. A PDF for the entire issue is labeled "print magazine" and is located near top right side of the IDU home page.

Table of Contents

(Click on page number to jump to articles)

Note: Intel does not control the content on other company's Web sites or endorse other companies supplying products or services. Any links that take you off of Intel's Web site are provided for your convenience.

Cover Story

New Chipsets Deliver Performance and Flexibility

David Puffer
Arizona Chipset Architecture Manager
Desktop Products Group
Intel Corporation

David Poisner
Component Architect
Desktop Products Group
Intel Corporation

The Internet and advanced applications are creating the opportunity to think in new ways, and Intel has risen to the challenge by designing a platform that meets today's business needs while allowing for compatibility with technology under development. The new Intel® Hub Architecture delivers the superior performance and high scalability required for the exciting e-Business and e-Home applications available today, as well as those in development for the future.

As the next step in the evolution of Intel® Hub Architecture, Intel recently introduced two new chipsets featuring enhancements designed for maximum performance and flexibility.

- *The Intel® 820E chipset* offers an innovative platform design that produces greater system responsiveness, reliability, and concurrency for overall performance gains.
- *The Intel® 815E chipset* is designed with integrated graphics and AGP4X plug-in feature for ease of upgrade across the mainstream market.

Both chipsets include an enhanced version of the I/O Controller Hub (ICH2) and deliver twice the I/O bandwidth over traditional PCI bridge architectures, while incorporating dedicated data paths to fully utilize the additional bandwidth.

Optimized Intel® Hub Architecture

The enhanced I/O Controller Hub (ICH2) is unique to the Intel 820E and 815E chipsets, and utilizes the latest I/O integration technology to provide cost savings and flexibility in system design and upgradability. There are three primary features in the ICH2 that distinguish it from the previous ICH architecture.

Support for ATA100—The ICH2 is designed to support the industry-standard ATA100 disk interface, increasing data transfer rates on IDE by 50 percent over ATA66 solutions. While ATA100 is ahead of most system capability today, recent increases in drive cache sizes, faster rotation rates, and increased platter density indicate that technology will soon make ATA100 a higher performance solution.

Integrated LAN Connect Interface (LCI)—The ICH2 enables three LAN options, with a single LAN driver for all three solutions. Working in conjunction with the integrated LCI, one external device supports both 10- and 100-Megabit Ethernet. An alternative external device connects to networks based on the Home PNA 1.0 specifications, and is compatible with Intel® AnyPoint™ networking adapters.

Additional USB Controller—By integrating a second USB controller, the ICH2 supports 24 Mbps bandwidth across four ports. This represents a significant increase over the 12-Mbps peak shared on systems with a single controller and an external 4-port hub.

Value-added Features

In addition to performance features in the ICH2, enhancements were made by Intel to enable soft audio/modem implementation, graphics flexibility, and innovative board configurations.

Six-channel AC '97 audio—While the original ICH offered stereo sound, the ICH2 supports either four- or six-channel digital surround sound. This enables software decode and playback of 5.1 and 6 channel audio formats, such as Dolby Digital (AC3), and DTS formats commonly found in DVD Video. Full surround sound can be used simultaneously with the soft modem connection to the Internet.

System manageability interface—Should system problems or failure occur, this system manageability bus slave interface allows external microcontrollers to access and read registers in the ICH2, and deliver solution commands such as resetting the chip or powering down the system.

Less interrupt sharing—An increase in the number of interrupt pins allows more peripherals on the system to operate without sharing their interrupts, thus increasing system performance.

Optimizes Performance

The Intel 820E chipset was designed to maximize the power of the fastest Intel® Pentium® III processors and deliver the performance that desktop users demand. This innovative platform design supports next-generation RDRAM*, doubling peak memory bandwidth from 800 Mbytes/second on PC100 SDRAM to 1.6-Gbytes/second peak, with numerous upgrade configurations. The RDRAM supports PC600, PC700, and PC800, and the more efficient high-speed 16-bit RDRAM interface enables greater board efficiency and memory granularity.

The Intel 820E chipset provides an AGP4X controller integrated in the Memory Controller Hub. This increase in graphics bandwidth allows the highest graphics performance and lifelike 3D, in addition to providing end-user upgraded graphics subsystems. The AGP interface with fast writes allows access to main memory at over 1 Gbyte/second, delivering the next level of graphics performance.

Flexibility

The Intel 815E chipset is designed to provide flexibility and ease of upgrade, while meeting the performance, stability, and reliability demands of the growing value and midrange desktop market. Compatible with both Intel Pentium III processors and Intel® Celeron™ processors, the 815E chipset offers ICH2 enhancements plus features that add value and are cost-effective, all without compromising system performance.

The Intel 815E chipset is the first production chipset to combine integrated graphics with AGP4X upgradability. Intel's 3D with Direct AGP is easily upgraded by adding a Graphics Performance Accelerator (GPA) card, AGP4X card, and/or a Communication and Networking Riser (CNR) card. The CNR card allows audio, modem, and/or LAN configurations on a single-base board design. This is also the first Intel® chipset to support PC133 SDRAM system memory, which significantly increases system performance compared to previous generation PC100 SDRAM chipsets. The flexibility of the chipset allows system developers to incorporate the 815E chipset into a single design that supports a wide range of desktop market segments.

Summary

The Intel 815E and 820E chipsets offer an ideal solution for bandwidth-hungry e-Business applications and complex computing environments. Smart I/O integration, dual USB controllers, enhanced audio and graphics, and support for future technologies are integral to achieving the performance, flexibility, and reliability that customers expect.

Intel's comprehensive efforts to enable the industry ensure fast deployment of next-generation platforms and a robust foundation to support emerging applications, with minimal risk.

More Info

For more information on the Integrated LAN solution, please read "ICH2 Enables Integrated LAN Solutions" in the Networking and Communications department of this issue of the *Intel Developer Update* magazine.

You can find in-depth information on the Intel 815E and 820E chipsets at the following areas of the Intel Developer Web site:

- Technical information on the 815E chipset
- Technical information on the 820E chipset
- Design guides for the 815E and 820E chipsets
- Datasheets for the 815E and 820E chipsets
- Drivers for the 815E and 820E chipsets
- Technical information on the Graphics Performance Accelerator Card
- Communication and Networking Riser (CNR) Card
- Technical information on the LAN
- Technical information on the AC '97

Author Bio

David Puffer has been with Intel for 12 years and is currently a chipset architecture manager in Arizona. During his first five years with Intel, David was a CAD tool software engineer, then switched to hardware and helped design the Intel® i960JX microprocessor and several products in the i960RX family of Intelligent I/O (I_2O) processors. After moving to the chipset group David became the lead microarchitect and handled much of the architecture for the Intel® 815 chipset. David has a B.S.E. in Computer Systems Engineering from Arizona State University, and was a 1996 Intel Achievement Award winner for development of the Intel® i960RP processor. David has one patent to his credit, with more in the works.

David Poisner is a component architect in Intel's Desktop Products Group (DPG) in Folsom, California. He has been with Intel for 14 years. During the past seven years, David has been an architect for various I/O components in the chipsets. He has been a contributor on various industry I/O standards and was primary author of the LPC specification. Prior to working in DPG, he was an applications engineer and manager in Intel's LAN controller group. David is a graduate of the University of Kansas with a degree in Electrical Engineering and has 14 patents to his credit.

Column

From the Editor

Donna Loveland
Managing Editor
Technology and Initiative Marketing
Intel Corporation

Column

In the northern half of the world, it's the time to get away from the desk and get out where it's hot. Well, this July, what's hot is on the desktop, and this issue *of Intel Developer Update* is here to tell you why. In addition, a selection of wide-ranging articles will bring you up to date on other news about Intel® technologies, platforms, and services for developers.

New Chipsets Deliver Performance and Flexibility, this month's cover story, focuses on Intel's innovative new 815E and 820E chipsets. Continuing the evolution of the Intel® Hub Architecture, these chipsets use the latest I/O technology to deliver performance, flexibility, and bandwidth to support complex Internet, e-Home, and e-Business applications.

ENERGY STAR and Instantly Available PCs is particularly timely, with the U.S. government's revised ENERGY STAR specs being released July 1. The article looks at a group of technologies that will be used to make the next generation of power-managed PCs not only smarter and faster, but also more economical and more environmentally friendly.

A Tool for Predicting a Novice PC User's Experience is an important outcome of the PC Ease of Use Roundtable. The Initial Experience Predictor (IEP) is a tool created to help OEMs improve consumers' initial ("out of box") experience with desktop PCs. This article covers the development, testing, and use of the IEP tool.

One Intel® Desktop Board D815EEA Fits Many Uses discusses the Intel Desktop Board D815EEA as a "one size fits all" performance PC solution. The board's innovative and scalable performance PC architecture not only supports cutting-edge technologies but also delivers the flexibility to meet a variety of price/performance targets.

Beyond-AGP4X describes an evolution of today's Accelerated Graphics Port (AGP) interface that can dramatically improve 3D graphics processing. Beyond-AGP4X can deliver increased bandwidth, hardware-enforced coherency, and other features to support the high levels of performance that games, animation, design automation, digital-content creation, and video editing will soon demand.

1394 Open HCI 1.1: What's New details the more critical enhancements and changes made to the 1394 Open HCI Specification Revision 1.0 to create Revision 1.1. This article gives an extremely abbreviated summary of the many differences between the two revisions and is intended for readers already conversant with the 1394 Open HCI Specification 1.0.

Universal Plug and Play Connects the Home walks through the five basic phases of UPnP, an architecture for peer-to-peer network connectivity for PCs, appliances, and other networked devices. UPnP offers consumers easy device and service connectivity through an open, standards-based architecture that utilizes TCP/IP and HTTP to enable seamless networking.

Intel® 82562 PLC Enables Chipset LAN Solutions discusses the next significant innovation in networking technology–integrating LAN solutions into a platform chipset–and how Intel achieved it with the new, second-generation I/O Controller Hub (ICH2) found in the new 815E and 820E chipsets, the companion 82562 Platform LAN Connect (PLC), and associated software.

Intel® 80303 I/O Processor Boosts PCI Bandwidth tells how Intel's third-generation I/O processor continues Intel's vision of meeting the growing demands of new Internet and storage applications. The 80303 enhances server I/O data flow and helps balance system-level performance by integrating a 64-bit, 66-MHz PCI-PCI bridge.

Intel® Developer Alliance Enables Better Solutions introduces a new area of the Intel developer site dedicated to software and Internet developers. The Intel Developer Alliance Web site provides the resources and tools developers need to create high-performance, cost-effective solutions for Intel® Architecture platforms–and bring them to market faster.

Whether you're reading this from your desktop or from a laptop on a sunny mountaintop, you'll find the height of developer information in July's *Intel Developer Update* magazine.

Enjoy.

Author Bio

Donna Loveland is the editor of *Intel Developer Update* magazine. She joined Intel's Platform Marketing group in 1999 as the editor of Platform Solutions News. Donna began her career with Intel in 1982 as a technical editor in an advanced microprocessor development group. Since then, she's held technical and marketing positions in leading-edge technology areas ranging from stereoscopic display to digital broadcast to scalable online content. Donna has a B.A. degree in English from the University of Rochester and an M.A. in Expository Writing from the University of Iowa.

Departments

Desktop

One Intel® Desktop Board D815EEA Fits Many Uses

Lisa Lau
Product Marketing Engineer
Reseller Products Division
Intel Corporation

David Magnoni
Senior Product Marketing Engineer
Intel Architecture Marketing Group
Intel Corporation

Overview

The Intel® Desktop Board D815EEA provides system builders with a "one size fits all" performance PC solution. This new board features an innovative and scalable performance PC architecture that not only supports cutting-edge technologies, but also delivers the flexibility to meet a variety of price/performance targets. Here's a quick feature-by-feature overview of this full-featured ATX motherboard:

Processors: via its PGA370 connector, supports Intel® Pentium® III and Celeron™ processors in FC-PGA package and Celeron processors in PPGA package. In addition, the board supports Pentium III processors with 100-MHz and 133-MHz Front Side Bus, and Celeron processors at 500 MHz and higher, with 66-MHz Front Side Bus.

Chipset: based on the Intel® 815E chipset, with DVI (Digital Visual Interface) support for flat-panel displays and digital CRT and TV-out capability.

Graphics: integrated Intel® 3D with Direct AGP. The board supports an optional Graphics Performance Accelerator (GPA) card in the AGP slot, enabling integrators to upgrade to a 4-Mbyte display cache.

Expandability: one AGP slot, four PCI slots, and one shared PCI/CNR (Communication and Networking Riser) slot. In addition, the board features two independent USB controllers, with four total ports, providing two USB rear connectors and a header for two front-panel USB connectors.

Connectivity: via the CNR, supports low-cost audio, modem, Home PNA* (Home Phoneline Networking Alliance), or LAN add-in cards. The board also supports an optional Intel® PRO/100 Network connection.

Memory: supports PC133 and PC100 SDRAM in three DIMMs, supporting up to 512 Mbytes of maximum memory.

Audio: options, including Creative Sound Blaster* PCI 128 (or ADI 1885 AC '97 Soft Audio).

BIOS: AMIBIOS and Intel® Rapid BIOS Boot.

Other Features: additional options, including a Digital Video Output header for a Flat Panel, Digital CRT or TV-Out card, and a Management ASIC. The board supports Instantly Available (Suspend to RAM) capability with Wake On USB and Wake On PS-2*. To simplify troubleshooting, the board features four diagnostic LEDs.

Processors and Chipset

The Intel Desktop Board D815EEA supports Intel's latest socket-370 Pentium III and Celeron processors and the highly integrated Intel 815E chipset.

To meet the requirements of the performance PC segment, the board supports Pentium III processors up to 933 MHz. For great quality with exceptional value, the board supports Celeron processors up to 600 MHz. For design flexibility, system stability, and performance headroom, the Intel 815E chipset is validated with Pentium and Celeron processors and provides a 133-MHz system bus, support for PC133 memory, AGP4X technology, and balanced bandwidth for memory and I/O.

The 815E chipset supports dramatic bandwidth improvements, including a 1-Gbyte/second memory interface and 1-Gbyte/second AGP bandwidth.

The chipset's integrated I/O Controller Hub (ICH2) provides two USB controllers to double the USB bandwidth available in the system. The presence of four USB ports eliminates the need for a USB hub down-on-motherboard. For high-performance systems, the ICH2 also supports the new Ultra ATA/100 disk drive interface while still maintaining backward compatibility with Ultra ATA/66 and Ultra ATA/33 modes. The ICH2 features an Integrated Network Media Access Controller for low-cost networking solutions and supports networking through the (CNR) interface.

CNR

The Communication and Networking Riser (CNR) provides system integrators with a flexible and low-cost way to implement audio, modem, HomePNA, or Ethernet technology. The CNR essentially adds networking technologies to the baseline functionality of an Audio Modem Riser (AMR). One of the major advantages of the CNR is the ability to share a PCI slot instead of taking up a dedicated PCI slot, as was the case with an AMR.

The CNR makes optimal use of the chipset's integrated media access controller. The CNR supports the chipset's LAN Interfaces. Intel® Platform LAN Connection (PLC) components enable networking functions. These include an Intel® Pro/100 Network connection for 10/100 Ethernet and an Intel® 1-Mbit/second Home PNA connection for basic home LAN. The Intel Pro/100 Network connection is ideal for SHO and consumer broadband networking and is compliant with the Wired for Management 2.0 and PC99 specifications.

The versatility of the CNR enables system integrators to use one motherboard for customers in both consumer and commercial markets.

The CNR supports full Plug and Play functionality.

Graphics and Video

The Intel Desktop Board D815EEA provides a graphics upgrade path for system designers in a single platform.

To minimize costs, integrators can design a system without a GPA card, using the integrated Intel 3D Graphics capabilities of the 815E chipset. Compared to Intel's 810E chipset, the integrated graphics engine in the 815E chipset uses Intel's proven drivers and adds tri-linear filtering and hardware acceleration.

- System designers can add a Graphics Performance Accelerator card in the board's AGP connector for improved graphics performance with 133-MHz local memory.
- Adding a third-party AGP graphics card enables scalable graphics performance.

For rich multimedia experiences, the board's Digital Video Output header supports digital video for DVI devices and TV output through add-in cards.

Audio

The current release version of the board features Creative Sound Blaster PCI 128 Audio, with 128-voice wavetable synthesis, 3D audio technology, and CD-quality 16-bit stereo digital audio performance.

Here is an at-a-glance summary of the hard audio capabilities:

- 128-voice polyphony and multi-timbral capability
- 16 MIDI channels
- 2-Mbyte, 4-Mbyte, and 8-Mbyte sample sets
- Support for Microsoft DirectSound* and DirectSound 3D* audio technology
- User-selectable algorithms for two or four speakers
- 16-bit stereo digital audio
- User-selectable sample rates
- Powerful amplifier
- Full-duplex support for simultaneous record and playback

Optional ADI 1885 AC '97 soft audio provides quality audio in an integrated solution. Features include 1,024 wavetable synthesis, 3D audio technology, and 16-bit stereo digital audio.

Easier Integration

Compatibility and platform reliability are key issues with any new PC board, and Intel works to ensure that all platform building blocks work together. The Desktop Board D815EEA and the 815E chipset were evaluated extensively in the development cycle. The result is a desktop board designed to work with other platform building blocks.

Intel Rapid BIOS Boot software technology makes using the PC more like using a consumer electronics device by providing dramatically faster access to PC applications from the time power is turned on. Rapid BIOS Boot results in significant reduction in Power ON Self Test (POST) time and the time required for the operating system to launch.

Intel® Express Installer 2.5 software suite, available with the Desktop Board D815EEA, helps shrink installation time. The software package includes Norton AntiVirus* software. Boxed boards come with Ultra ATA/100, DMA/33, and floppy disk drive cables, I/O shield, and product documentation.

The Desktop Board D815EEA is available in two SKUs, offering a choice of CNR or integrated LAN. Both SKUs feature Digital Video Output headers and diagnostic LEDs.

Summary

The Intel Desktop Board D815EEA enables system integrators to stock one performance PC motherboard to meet a variety of user profiles and price points ranging from business PCs with integrated networking to systems with the screaming performance to support the latest 3D graphics and games.

The board provides flexibility for system builders and the ideal migration path for current Intel® SE440BX-2 customers. The Intel Desktop Board D815EEA provides a validated foundation for performance desktop platforms based on Intel Pentium III and Celeron processors and the 815E chipset, backed by Intel quality, technology, and support.

With its rich feature support and tremendous versatility, this desktop board is a versatile integration solution, "one board with many uses."

More Info

Information on Intel® Boxed Desktop Boards is available on the Intel Channel Web site.
Also see the Desktop Board pages provided on the Intel Developer Web site and Intel's Support Web site.

Author Bios

Lisa Lau joined Intel in 1992. As a product marketing engineer in the Reseller Products Division, Lisa has worked on a number of projects, including the motherboard designed to support Intel's Willamette processor. She holds a B.S.E.E. from San Francisco State and an M.B.A. from the University of New Mexico.

David Magnoni joined Intel's OEM Platform Solutions Division in 1999 as a product marketing engineer. He has seven years of industry experience in marketing, information technology, and consulting. A Hertz Foundation Fellow, Dave holds degrees in systems engineering from the U.S. Naval Academy and in computer science from New York University. His M.B.A. comes from the MIT Sloan School of Management.

Beyond-AGP4X

Ed Schmit
Workstation Marketing Initiatives Manager
Intel Architecture Marketing Group
Intel Corporation

Anne Gregory
Visual PC Initiative Manager
Intel Architecture Marketing Group
Intel Corporation

Overview

Accelerated Graphics Port (AGP) is an interface that dramatically improves the processing of three-dimensional (3D) graphics. However, today's AGP interfaces will soon have difficulty meeting the high levels of performance that will be required by upcoming applications. These applications include games, 3D graphics animation, mechanical design automation, digital-content creation, and video editing.

One issue for handling these applications is increased bandwidth for data traffic. Another, more complex issue is the potential need for isochronous mode operations. To address these issues, initiatives are already in process for a more advanced form of AGP. In particular, Intel is working on an evolutionary initiative called Beyond-AGP4X.

Evolutionary Design

The Beyond-AGP4X initiative is an evolutionary approach to the next level of efficiency in graphics traffic. Evolutionary designs for improving the graphics attach-point are the most cost-effective approach. By supporting both existing and future protocol and interconnect requirements, Beyond-AGP4X cards can be made compatible with existing AGP4X boards. In addition, motherboards can easily be designed to support both AGP and Beyond-AGP technologies that continue to run AGP-compliant graphics cards. This backward compatibility reduces the hardware costs for upgrading a system for greater graphics performance.

Beyond-AGP4X offers several advantages over AGP4X. Beyond-AGP:

- Is an evolutionary design based on APG2.2 protocols, which maintains the established AGP interface and protocols.
- Maintains connector-level compatibility with AGP. Board impedance and routing constraints will also be similar to AGP.
- Uses a new, more efficient signaling scheme to support the operation speeds required for doubling (or for even greater increases in) the bandwidth of AGP4X.
- Can be implemented with minimal interruption to the graphics industry. This is because of the evolutionary nature of the interface—the backward compatibility of Beyond-AGP4X with the existing AGP software infrastructure.

To ensure widespread adoption of Beyond-AGP4X, the interface cost structure will be similar to AGP. The Beyond-AGP4X specification will also maintain AGP and AGP Pro power-delivery schemes. The overall industry transition from AGP4X (or AGP2X) to Beyond-AGP4X will be straightforward.

Dataflow Models

There are two main types of workstation-class graphics applications that really push the limits of the graphics subsystem: mechanical design automation and digital content creation. For each of these application segments, the graphics subsystem usage models are somewhat different. Each usage model determines the flow of graphics data through the platform. In turn, the unconstrained flow of this data determines the overall achievable graphics performance.

The interface between the graphics subsystem and the rest of the platform is only one link in a more complicated dataflow infrastructure. This infrastructure must balance all usage models in order to effectively exploit all graphics capabilities of the system. Simply increasing the bandwidth of the graphics interface is not likely to produce the desired increase in graphics performance. Instead, increasing bandwidth alone will create dataflow bottlenecks in other interfaces, such as those to the main memory subsystem or the processor(s).

An Example

For example, in one dataflow model, the graphics database resides in main memory. Based on the application's needs, the processor reads the database, modifies the database in some fashion, and writes the database back to a buffer in main memory. Following this action, the processor directs the graphics subsystem to fetch the modified data from the main memory buffer. The graphics subsystem then renders the modified data to a local frame buffer for eventual display.

When using the OpenGL* (graphics libraries) immediate-mode API for this data manipulation, the processor transfers a complete, new frame of graphics data each time the sequence is executed. This complete, new frame of graphics data is transferred to the graphics subsystem, resulting in a sustained dataflow pattern through the platform.

The sustained dataflow pattern effectively requires that, for each byte of graphics data transferred to the graphics subsystem, 2 bytes will be transferred on the processor's system interface (front-side bus); 3 bytes will be transferred on the interface to the memory subsystem. If AGP4X is to be fully utilized at 1 Gbyte/second, the implication of this dataflow model is that the platform must support 2 Gbyte/second on the front-side bus and 3 Gbyte/second on the memory interface. (Other usage models result in different dataflow requirements with different platform implications.)

Balancing Usage Models

The previous example shows why designers must understand the target application environment, and why they should ensure a balanced platform infrastructure. Ensuring a balanced platform infrastructure is the best way to make sure that graphics performance expectations are met.

For example, consider a digital content creation application that consists of 3D animation with a high-definition video stream as texture. These applications typically use the immediate-mode OpenGL API for communicating with the graphics subsystem. Even today, a typical professional 3D application could use one million polygons per frame. Assuming an average of 1.2 vertices per polygon and a real-time frame refresh rate of 30 fps, the result is 36 million vertices per second to the graphics subsystem.

The size of a vertex of data depends on the graphics subsystem and such variables as where the 3D geometric transformation and lighting calculations are performed. If you assume an average of 32 bytes per vertex, the 3D graphics traffic alone would be around 1.1 Gbyte/second. In addition, the high-definition video stream used as texture would result in approximately 230 Mbytes/second. The total sustained traffic on the graphics interface could be around 1.33 Gbytes/second, which clearly exceeds the capability of AGP4X.

With Beyond-AGP4X, bandwidth capacity is double that of AGP4X—and is more than capable of handling the amount of traffic required in this example. In addition, Beyond-AGP4X will be able to support increased bandwidth in the future.

Enhancements

Beyond-AGP4X offers several performance and feature enhancements, as well as some simplifications of the AGP specification. Some of the enhancements are optional and targeted at the higher end of the workstation market.

Significant changes include:

- Twice the AGP4X bandwidth, achieved through the use of a new signaling scheme.
- Relaxation of ordering rules in order to more efficiently process graphics system traffic to and from system memory.
- Limiting the transaction size of data to the graphics subsystem in order to enable efficient use of the interface.
- Simplification of the interface usage model, which includes removing the notion of high-priority transactions. The usage model simplification also includes using an alternate, lower performance mode of transaction request generation (called PIPE) by the graphics subsystem.
- Providing a defined mechanism to enable high-performance cache coherent transactions on the interface.
- Specifying an optional infrastructure that supports isochronous data streaming on the interface.

Isochronous Mode Operations

Certain applications, specifically in the digital content creation area, require a predetermined and timely flow of data. In such applications, any interruption in the timely flow of data will result in errors in the way the data is processed.

An example of this would be a stream of video intended for display in real time. Such a stream would require dataflow at the frame rate dictated by the video monitor. For example, a standard television (TV) video stream in red-green-blue color space would need a frame rate of approximately 500 KB every 33 ms. (In other words, the transfer of each complete frame must be completed within 33 ms in order to create a smooth, real-time video stream.) A video card processing this stream may be able to buffer an entire incoming frame while working on a previously received frame. If the frame does not arrive within the allocated time, the video card may have to skip the entire frame, resulting in a flaw in the display.

The Isochronous Contract

The real-time buffering and processing of graphics traffic described above is called isochronous traffic. Unlike Beyond-AGP4X, the AGP interface does not provide a mechanism for isochronous traffic. Instead, graphics traffic in AGP is essentially asynchronous. In contrast, Beyond-AGP4X provides an optional infrastructure for enabling the isochronous flow of data in the midst of the existing asynchronous flow.

For graphics to work properly with isochronous-mode operations, the video card establishes an isochronous contract with the workstation or PC platform. The contract is based on the capabilities of the card and the platform.

A typical usage model where isochronous and asynchronous traffic may coexist would be one in which the graphics subsystem integrates video capabilities, and the application mixes the video with 3D graphics traffic. In the standard, real-time TV video stream, for example, the isochronous contract could be delivering 166 KB to the video card every 10 ms. With a transfer rate of 166 Kbytes/10 ms, the isochronous contract achieves the same throughput requirement of asynchronous traffic, and also accommodates the capabilities of the existing graphics card and a given platform. With isochronous traffic, the platform can vary the dataflow throughput within the 10-ms period (the isochronous period). The allowed variability within the isochronous period is referred to as jitter.

One of the issues in isochronous-mode operations is that the isochronous traffic must ensure that the entire graphics frame is transferred successfully. The Beyond-AGP4X specification will describe how this verification is achieved.

Isochronous Specifications

The isochronous capability in Beyond-AGP4X is based on a contract established between the card and the platform. With Beyond-AGP4X, a new set of transactions is created to enable the card to perform isochronous transaction requests to main memory.

The Beyond-AGP4X specification will fix the isochronous period. The specification will also fix the maximum latency between a request from the card and the platform's response. Additional mechanisms will establish the number of isochronous transaction requests that the card can initiate each period. To ensure that the platform can meet isochronous requirements, the Beyond-AGP4X specification will also:

- fix the size of the data packets
- define new, relaxed ordering requirements
- remove link-level flow control
- remove any hardware-maintained coherency requirements that pertain to processor caches

As already indicated, this Beyond-AGP4X feature is optional and is targeted specifically at higher-end platforms.

More Info

A presentation on the Beyond-AGP4X interface can be found on the Beyond-AGP4X Web site (pdf 1584 KB). In addition, you can read about 3D graphics issues, expected application requirements, and possible approaches for bandwidth and timing solutions in the April 27, 2000, *Electronic Design News* (EDN) article titled "Balancing in Three Dimensions."

Summary

Workstation and desktop applications are continually scaling and pushing the limits of the graphics subsystem and the graphics attach-point. The key segments in workstations that drive 3D graphics performance are in the mechanical CAD and digital content creation segments. The advent of high-definition video streaming and its integration into the graphics subsystem pose new challenges to the graphics interface.

On the desktop, 3D graphics technology is used primarily to make realistic PC games. Consumer applications, such as these realistic PC games, are a key driving force for moving the bus beyond AGP4X on the desktop. As the demand for creative and compelling Internet experiences increases, 3D technology will be used across more Web-based applications.

To address the issues presented in this article, work has already begun on an evolutionary approach called Beyond-APG4X. This protocol will increase bandwidth, make data transmissions more efficient, and provide means for compatibility with existing AGP hardware and software infrastructures. In addition, the timing of the introduction of the new, Beyond-AGP4X graphics interface will also coincide with increased capabilities in Intel® Architecture workstation and PC platforms.

Beyond-AGP4X will make it easier for users to upgrade to newer technologies. It is a solution that can be implemented with minimal interruption to industry, and which will meet the needs already being anticipated for high-end graphics processing.

Author Bios

Ed Schmit is the workstation marketing initiatives manager in Intel's Architecture Marketing Group. He has been with the company for four years, and has worked on a variety of initiatives, most recently with graphics initiatives. Ed holds a B.S. in Materials Science from MIT, an M.S. in Mechanical Engineering from the University of Washington, and an M.B.A. and M.S. in Manufacturing from the University of Michigan.

Anne Gregory has been with Intel for four years, and is currently the visual PC initiative manager for 3D and DVD technologies in the Architecture Marketing Group. She holds a B.S. in Pure and Applied Physics from the University of Manchester Institute of Science and Engineering, United Kingdom, and an M.S. in Electrical Engineering from Santa Clara University.

ENERGY STAR and Instantly Available PCs

Jill Abelson
Communications Director
ENERGY STAR Program
U.S. Environmental Protection Agency

Overview

"In the future, computers will weigh less than 1.5 tons." Popular Mechanics, 1949.

This was almost an understatement 50 years ago, and it told only half the story: thanks to advanced technologies, computers have not only shed a few pounds, they've become a ubiquitous part of life on this planet. And given their proliferation—100+ million PCs in the U.S. alone—computers are also having an effect on the environment.

To address this issue, the EPA launched ENERGY STAR (E*), a voluntary partnership with equipment manufacturers across several industries. Equipment with the ENERGY STAR label reduces power consumption and the pollution associated with the use of electricity by powering down and going to sleep during periods of inactivity.

Since demonstrating the first ENERGY STAR-qualified PC in 1994, Intel has continued to develop advanced power management technologies. The latest of these, Instantly Available (IAPC), is a group of technologies that will be used to build the next generation of deeply power-managed PCs.

Through better, advanced technologies like IAPC, Intel is making computers not only smarter and faster, but also more economical and more environmentally friendly.

Conserves Energy

Energy use by office equipment is one of the fastest-growing sources of electricity consumption in businesses and homes. It currently accounts for more than seven percent of total commercial sector electricity use. Much of this energy is wasted because office equipment sits idle for long periods. During the average workday, PCs are actively used for only about 4 hours, with another 5.5 hours being idle time. In some environments, PCs are left on 24 hours a day, creating even more downtime and added energy use.

To tackle this problem, the EPA launched ENERGY STAR, a voluntary partnership with office equipment manufacturers and other manufacturers across several industries. The ENERGY STAR label identifies energy efficient products that save money by eliminating wasted energy.

Reduces Pollution

Office equipment with the ENERGY STAR label saves energy by powering down and going to sleep during periods of inactivity, but still offers all the performance features of standard equipment. On a monthly or annual basis, this equipment uses about *half* as much energy as standard equipment, saving its owners millions of dollars in electricity costs. The EPA is interested in power management for another important reason: in office equipment, power management offers huge potential for dramatically cutting air pollution associated with electricity use.

IAPC Benefits Users/Environment

Since demonstrating the first ENERGY STAR qualified PC in 1994, Intel has continued to develop advanced power management technologies that reduce total PC power consumption. Traditional PC power management techniques, while broadly successful, do not address user requirements: resume times are not acceptable, network connections are often broken, and systems are noisy when asleep. And while power management is a feature of many PCs, many have had their power management functions disabled. Intel's current efforts are aimed at making power management both easy to use and environmentally effective.

Instantly Available (IAPC) is Intel's term to describe a group of technologies that will be used to build the next generation of deeply power managed PCs. These PCs can power down into a very deep sleep state, and yet wake up fully within five seconds while still retaining the capability to respond to external or user-programmed events. IAPC is an open, license-free architecture based on open industry standards.

IAPC technology offers major benefits for both business and home users:
Maximizes energy cost savings and reduces total cost of ownership for business owners.
According to EPA estimates, a medium-sized business with 1,000 PCs would save $171,000 (US) per year by deploying PCs with this technology instead of PCs without power management capabilities.

Enables powerful management tools that can remotely wake up PCs at night and update software *configurations.* Employees experience no interruptions; support is easier for IT managers.

Enables out-of-hours data/information gathering, keeping workers better informed and more *productive.* IAPC allows PCs to host personal web pages, share data, run smart software agents, and analyze data, all before going back to sleep.

Provides new applications for home PCs, including home networking and messaging. Virtually eliminates the "boot" process, waking up instantly and returning to any previous activities. For home users, PCs are always connected and immediately available.

IAPC Far Exceeds New e*

The EPA has revised current computer specifications, making them more stringent for two reasons: to encourage additional improvement in power management, and to pursue greater pollution reduction goals.

Knowing that many products on the market are already compliant and that better power management technologies were becoming available, the EPA believed the time was appropriate to raise ENERGY STAR standards. For higher end workstations, a different set of power specifications will apply, based on power supply size. These systems typically use more power than others on the market; thus, the new specifications encourage their use of power management as well.

Instantly Available actually exceeds new ENERGY STAR specifications for computers. The EPA's new specifications, which become effective July 1, 2000, reduce the sleep level from 30 watts/sleep mode to 15 watts/sleep mode for most standard computers. With IAPC technology, PC manufacturers can design products that go to sleep and consume less than 5 watts—*a sixfold improvement over sleep states required by ENERGY STAR.* IAPC technology allows even greater energy, cost, and pollution savings.

EPA View on Benefits

From an EPA perspective, Intel's technology advancements can offer tremendous environmental benefits. First, Instantly Available technologies will contribute to reduced global energy consumption. According to EPA estimates, PCs with Intel's IAPC technology would consume 63 percent less energy per year than non power managed PCs, saving the U.S. alone *$9 billion* in energy costs over the next 10 years. Furthermore, these advancements will help reduce the air pollution that contributes to global warming and other health and environmental problems. Instantly Available technologies have the potential to reduce thousands of tons of carbon dioxide pollution annually, equivalent to removing 18 million cars from the road each year.

Because IAPC technology is so easy to use it will encourage wider adoption of power management. With Instantly Available, users will be less likely to disable power management features. The EPA hopes that this "ease of use," as Intel calls it, will help increase power management compliance and enabling rates, leading to less power consumption overall. Even when PC users leave their computers running, more efficient equipment will permit both energy and pollution savings.

More Performance

The key feature of Intel's technology*: it achieves sleep state power consumption with no compromise to system performance or usage.* In fact, IAPC adds capabilities to PCs. The technology removes any conflict between the network and reducing energy costs. That alone makes power management more attractive and ultimately more widely used.

Consumer Awareness

Since launch of ENERGY STAR Office Equipment, the EPA has worked consistently to educate OEMs, end users, and consumers about the benefits of power management. It hasn't been easy. The EPA knew as many as 50 percent of all qualified PCs in business use have their power management disabled. Now, the technology is there to make those features seamless and easy to deploy. OEMs like Dell Computer, Compaq Computer, IBM, Inc., Ricoh Corp., and Canon USA, Inc. have supported EPA initiatives and have worked to educate end users, dealers, and the public. For a list of all ENERGY STAR Office Equipment partners, visit the ENERGY STAR Web site.

Consumer electronics like TVs and VCRs have a similar power management problem: they waste energy in stand-by mode, when we think they are actually turned off. Interestingly, this market segment views power consumption as an issue; it has captured the attention of consumers and the media alike. ENERGY STAR partners like Panasonic (Matsushita Electric Corporation of America), Sony Corporation of America, Aiwa America Inc., Sharp Electronics Corporation, Toshiba America Inc., and Memorex are taking advantage of their products' energy efficiency to gain a market advantage over other manufacturers who do not incorporate energy saving features. These partners, like Intel, set themselves apart from others in industry through their commitment to the design, manufacture, and marketing of environmentally responsible products and technologies.

Summary

Office equipment is one of the fastest-growing sources of electricity consumption in businesses and homes. PCs, which are actively used for only a fraction of the time they're turned on, account for vast amounts of wasted electricity. This idle time is expensive not only for the owners of these PCs but also for the environment, due to the air pollution associated with the use of electricity.

To address this problem, the EPA launched ENERGY STAR, a voluntary partnership with manufacturers across several industries. Office equipment with the ENERGY STAR label saves energy by powering down and going to sleep during periods of inactivity, but still offers all the performance features of standard equipment.

Intel's Instantly Available (IAPC), a license-free architecture based on open industry standards, exceeds the EPA's new ENERGY STAR specifications for PCs, which become effective July 1, 2000. From an EPA perspective, Intel's technology advancements can offer tremendous energy savings. Because IAPC technology is so easy to use, it will also encourage wider adoption of power management. What's more, IAPC opens the way to new applications.

EPA hopes that developers not only continue to improve power management technologies, but that they also strive to use it in their own business environments.

More Info

Information about Instantly Available technology and Power Management can be found at the Instantly Available Technology Web site. For more information on ENERGY STAR, go to the ENERGY STAR home page.

See and hear about Intel's latest updates of Instantly Available technology efforts at the Fall '00 Intel Developer Forum Conference, August 22 through 24, in San Jose, California. Details and registration information is available at the Conference Web site.

If you are interested in PC consumer promotional opportunities around ENERGY STAR and Instantly Available, contact Nancy Sumrall, Instantly Available Program Manager,
Technology Initiatives Marketing, Intel Architecture Marketing Group at (480) 554-1703 or by email at nancy.l.sumrall@intel.com.

Author Bio

Jill Abelson is a communications director with the U.S. Environmental Protection Agency's ENERGY STAR Program. She manages public education, media outreach, and corporate relations for ENERGY STAR products, and works closely with a number of ENERGY STAR partners on collaborative communications and marketing efforts. Prior to joining the U.S. EPA in 1993, Jill worked for Florida's Department of Environmental Regulation, several environmental non-profit groups, and the U.S. Senate. Her writing on global climate change has been published in newspapers across the U.S. She holds a B.A. from Mount Holyoke College.

A Tool for Predicting a Novice PC User's Experience

Paul F. Sorenson
Manager, User Centered Design Group
Intel Architecture Labs
Intel Corporation

Overview

The PC Ease of Use Roundtable was started in August of 1998 as an informal, cross-industry body made up of companies that could work together to improve PC Ease of Use. The objective of the Roundtable is "To identify and drive awareness of the ease of use of consumer PCs through defining desired user experiences for key tasks and the opportunities to improve it."

Many major OEMs involved with PCs, along with other companies, have been attending and collaborating on a series of initiatives focused primarily on extending the growing base of new PC users who benefit from improving ease of use. As part of this work, all members shared data about highest impact ease of use problems they were seeing with their products. A key area they identified for their attention was the initial experience and setup of the PC.

The Roundtable developed the Initial Experience Predictor (IEP) as a tool to help OEMs improve consumer PC users' initial ("out of box") experience with consumer desktop PCs. This article covers the development, testing, and use of the IEP tool.

Development and Testing of the IEP Tool

The flowchart in Figure 1 shows how the PC Ease of Use Roundtable developed its Initial Experience Predictor tool.

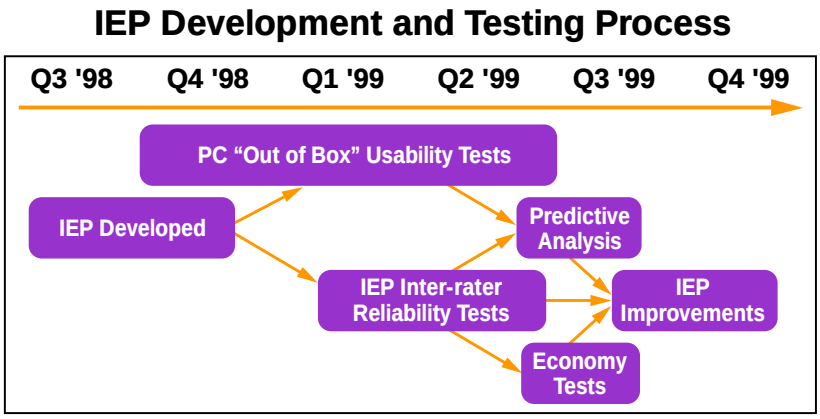

Figure 1

Selecting the original questions—An initial list of items was selected from a series of tools already available, including QUIS (Questionnaire for User Interaction Satisfaction), internal Intel guidelines, work by NIST (National Institute of Standards and Technology), and other sources. In addition, customer support organizations from several computer manufacturers provided information about issues they face in supporting novice users' out-of-box experience. All of this input was incorporated into the initial draft.

Refining based on usability tests—A usability test scenario was developed that mapped user-tasks to the items in the list. Two rounds of usability tests were conducted in Q4 '98 and Q1 '99 with a total of 58 participants using 9 systems (5 to 8 participants per system). The participants were all Portland-area residents who did not own a PC at home.

Each usability test participant began the session with a system fully packaged as it would be sold in the retail channel. The participant set up the system, turned on power to all components, connected to the Internet, registered the product online, and performed key first-use tasks such as sending e-mail, writing a letter, and listening to a music CD.

Results from this testing were used to remove items that dealt with PC characteristics that did not strongly influence how well users performed in unpacking, setup, and initial use of these systems. Characteristics that were not on the original list but did have strong influence in testing were added. The usability tests were also used to derive a preliminary sense of how important each item is (a simple "high, medium, low" designation). From this, a scoring system was developed so that raters could assign point values to individual characteristics.

To further refine the checklist items, usability experts from IBM Corporation conducted two rounds of tests in Q2 '99. In these tests, the experts set up a number of different systems while rating the systems against version 0.6 of the checklist. They captured comments and suggestions for revising the checklist during the setup experience. Versions 0.7 and higher of the checklist have incorporated these comments.

Testing IEP Validity

The goal of this project was to build a tool that would be:

- *Reliable*—Scores for a given system remain constant even among different evaluators.
- *Valid*—Scores for a given system correlate highly with performance measures in usability testing for the same system with novice users.
- *Economical*—Using the tool adds no more than 30 minutes to the process of setting up a new PC system.

The Roundtable tested these goals in a series of studies, the results of which are summarized below:

Reliable

In one round of testing conducted at Intel in Q2 '99, 48 expert evaluators set up 7 different PCs (5 to 7 evaluators per system, working alone or in pairs). During the setup, they rated the system on version 0.6 of the checklist.

From their results, we assessed reliability in two ways:

- We measured the correlation of the scores *across all questions* for individual raters of a particular system.
- We looked at the stability of scores *for each individual question*. We calculated the coefficient of variation (standard deviation divided by average) of each question for each system. If a question had a coefficient of variation greater than 0.5 for three or more systems, we selected the item for revision.

Result: For all systems, raters' scores were highly correlated ($r^2 > .90$). However, based on the variance among scores on the same system, we needed to revise 20 percent of items in this draft of the checklist.

Valid

We assessed validity by comparing results of the 48 evaluators in the Q2 '99 Intel study to the 58 end-user usability tests conducted over Q4 '98 and Q1 '99. We compared the scores that evaluators gave in each section of the checklist to the average completion times and intervention rates end users required to complete those tasks.

In this analysis, we defined predictive in two ways—one quantitative, one qualitative.

Quantitative: We defined a question as reasonably predictive if:

- The ratings were correlated *in the desired direction* with all three end-user performance measures. For example, ratings should correlate with "time to complete" in a negative direction (ratings increase as time to complete decreases). Ratings should have negative correlation with interventions, and positive correlation with success rate.
- The correlation coefficient (Pearson's r) was 0.3 or greater for at least two of the end-user performance measures.

Qualitative: Based on our observations of end users in earlier usability testing, we again categorized the IEP questions by importance to the success of the end-user experience (in a simple high, medium, or low rating).

Result: 12 percent of the questions met all of the quantitative predictability requirements; the rest did not. The qualitative analysis showed that 42 percent of the questions were critical to end-user success. Based on these two results, we dropped 58 percent of the questions from the earlier version. For the remaining 42 percent, we assigned high score values to the questions with good predictive power.

Economical
In a round of baseline usability tests with expert evaluators, the average time to set up a brand new system out of the box was just over 45 minutes. In a second round in which evaluators also answered questions on version 0.6 of the checklist, the average completion time was more than 96 minutes. This duration may even be artificially low, because after about 90 minutes, some evaluators were unwilling to spend any more time answering the questions and many skipped the last section.

This resulted in revision 0.7 of the checklist, with a reduced set of questions and presented as a software application. In a final round of usability testing with expert evaluators, the average completion time was just over 66 minutes, meaning it added about 21 minutes to the baseline set-up times. This indicated that the final software version of the checklist met our goal of adding no more than 30 minutes to the process.

How to Use the IEP Tool

The Initial Experience Predictor (IEP) tool was designed as an online utility. There is a way to print out the text of each question, but the printed version does not provide the flexibility of the software utility. It is highly recommended that you use the online version if at all possible. If you plan to use the hard-copy version, the Roundtable recommends you read through the questions before you start any evaluation.

This is the process you should follow to conduct your own testing:

1) Install the IEP program on a laptop or other system that you can keep nearby during the new PC setup and evaluation.
2) Choose the system to be evaluated. To take full advantage of the IEP, the evaluation should start while the system is in its packaging, including documentation and accompanying software.
3) Before unpacking the system, launch the IEP application and start a new checklist (File, New Checklist). Enter your name and a session name that will help you identify this particular evaluation later on.
4) Look through the questions in the checklist to get an idea of the kinds of questions you will encounter, and then answer the questions while you are unpacking and setting up the system. The checklist contains four sections:

Section 1, Hardware Setup: Answer the questions in this section as you open the packing boxes, put the system hardware together, and turn on the power.

Section 2, Software Configuration: Answer these questions as you step through the first-time OS configuration and driver installation.

Section 3, Internet Connection: Answer these questions as you set dialing properties and establish an ISP account for the first time.

Section 4, Initial Software Discovery: Answer these questions by exploring the options available for e-mail, Web browsing, and basic office applications.

5) After completing the checklist, choose View Score to see how the system rates. (You can view partial scores at any time while completing the checklist.) The View Score screen contains three types of numbers:

Actual score: The number of points accumulated in the current session.

Possible score: The maximum number of points possible, either for a particular section or for the entire checklist.

Percentage of possible: The percentage of total possible points the actual score represents.
The View Score screen will show these numbers for all four sections and for the entire checklist. If you have not completed the checklist before viewing the scores, the "completed questions" entry will be less than 100 percent.

6) View, save, or print a report of the results by choosing Current Checklist Results under the View Reports menu. You can view a formatted report of all the questions and answers for the current session.

7) To compare scores among different systems or different evaluators, choose Text File Summary By [Evaluator, System, etc.] under the View Reports menu. This generates a space-delimited text file containing the results of multiple sessions sorted by evaluator, date, session name, or system. This file can be imported into an analysis package (e.g., Microsoft Excel*) for further data processing.

What the Scores Mean

The final scores generated by the IEP will help predict how successful a very novice end user will be in setting up the system. Each question in the IEP has an associated score value based on how important that question is in predicting the end user's success. Very important questions can change a system's score by as much as 50 points. The reports under the View Reports menu show the point values for individual questions.

To improve the score for a particular system, view the Current Checklist Results report to see which section or questions suffered the biggest point losses. Then refer to the Design Guidelines that accompany IEP (available at the Ease of Use Roundtable Web site). The guidelines offer specific suggestions for improving scores in all four sections of the checklist.

Summary

The PC Ease of Use Roundtable, an informal industry body, has been working since 1998 to make PC use more enjoyable for consumers and less costly for vendors. In pursuing its charter to identify and drive ease of use, the Roundtable has created a number of opportunities for developers to improve PC products. The findings and recommendations documented in its numerous white papers and tools have already come to market in products that make PCs easier to use. The Roundtable makes its materials available to the developers free of charges and fees, and it encourages broader industry participation in its ongoing efforts.

The Roundtable meets monthly, and its participant list has grown to include PC OEMs, communications, networking, and peripheral providers. Participants include Intel Corporation, 3Com Corporation, Altec Lansing Technologies, Inc., America Online, Inc., Aveo, Inc., Canon, Inc., Cisco Systems, Compaq Computer Corp., Creative Technology, Ltd., Dell Computer Corporation, Gateway, Inc., the Hewlett-Packard Company, NEC, Inc., IBM Corporation, InFocus Systems, Inc., Iomega Corp., Eastman Kodak Company, Lexmark International, Inc., Logitech, LucentTechnologies Inc. Virtual Networks, Nortel Networks, and Samsung Electronics Corporation.

More Info

The Ease of Use Roundtable Web site features an IEP area with the links to the IEP checklist and an automated version of the checklist as well as links to additional information including IEP collateral, background on IEP, and IEP design guidelines.

Author Bio

Paul Sorenson manages Intel's User Centered Design Group, an internal consulting team focused on understanding how end users interact with technology. He began his career at Intel in 1994 with the company's PC Enhancement Operation, which marketed Intel's first consumer products. Paul has been a core team member on the PC Ease of Use Roundtable since its inception. He leads Roundtable efforts on the Design for Supportability Guidelines and the Initial Experience Predictor Tool now in use by Intel and many PC OEMs. Before joining Intel, Paul spent more than a decade at IBM Corp., Lockheed Martin Corp., and the Hewlett-Packard Co. working on a wide range of products, including the first PC mice, advanced User Interface prototyping systems, and the first palm-top PCs. Paul holds a bachelor's degree in biology from Willamette University and a master’s degree in Neurophysiology/Linguistic-semantics/Experimental Psychology from the University of Oregon. His doctoral work was with the University of Texas, Austin, in Human Experimental Psychology.

Initiatives and Technologies

1394 Open HCI 1.1: What's New

Steve Bard
Manager, Mobile Interconnect Group
Intel Mobile Technologies Lab
Intel Corporation

John Fuller
Program Manager, External Buses
Core OS Device Drivers Gp/MS Windows*
Microsoft Corporation

Overview

First published by the 1394 Open HCI Promoters Group in October 1997 as Version 1.0, the 1394 Open Host Controller Interface (Open HCI) Revision 1.1 is the next instantiation of an implementation of the link layer protocol of the 1394 Standard for a High-Performance Serial Bus. Revision 1.1 incorporates a number of improvements, fixes, and new features as well as further articulation of Revision 1.0 features and functions. This article gives an extremely abbreviated summary of changes to the specification. It is intended for readers already conversant with the 1394 Open HCI Specification 1.0.

New and changed content contained in Revision 1.1 has been categorized into three classifications: critical, important, and significant but less important. This article provides a brief description of the subject areas in the first two classifications. Those not discussed in this article have been determined to be easily understood by reading about them in the specification.

Critical classification subjects include the following:

- Isochronous Receive
- Ack_data_error
- CSRcontrol register
- New effects of busReset Event
- Power Management and ack_tardy
- Out of Order Pipelining

Important classification subjects include the following:

- RegAccessFail
- AT PHY Packet Transmit
- Changes to ITDMA
- Changes to Autonomous CSR Resources
- Physical RequestFilter Registers
- SelfID Changes
- tag1syncfilter

A significant, but less important, classification subject is Mini ROM.

Critical Classification

Isochronous Receive

Multi-channel Mode (Section 10.3.2)

Because the multiChanMode bit is undefined after reset, software should initialize it in all IR contexts (even inactive ones).

Dual Buffer Mode (Sections 10.1.2 & 10.2.3)
Dual Buffer Mode is a new isochronous Receive DMA mode that allows the 1394 header, time stamp, and the first quadlets of data to be placed in a different buffer than the rest of the data. The rest of the data is "compacted" for presentation to a decoder. It works when incoming packets vary in size unpredictably.

Why do Dual Buffer Mode? Sometimes software needs to strip 1394 and transport protocol headers from the real data to create a data-only buffer. If this is done by software moving the data, it uses additional memory bandwidth—a significant amount for some audio and most video formats. Figure 1 shows the formant of a new DMA descriptor that describes the two buffers used for Dual Buffer Mode.

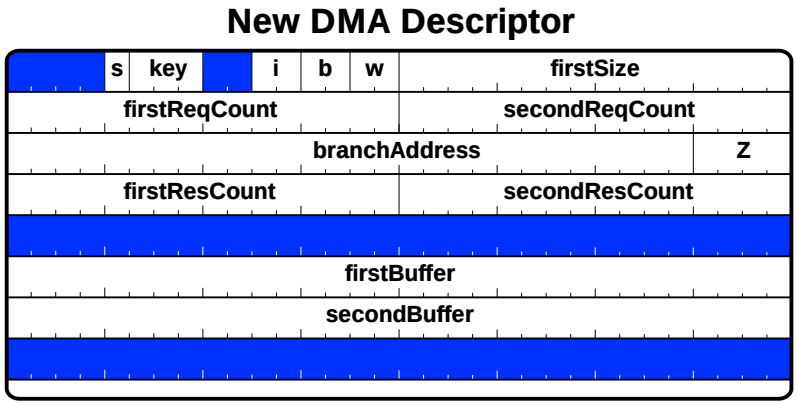

Figure 1

In Dual Buffer Mode, the 1394 packet header and the stream's protocol header, which are fixed length, are stripped and placed in a buffer separate from the buffer containing actual data transported by the stream. Figure 2 shows how this works for two successive isochronous packets.

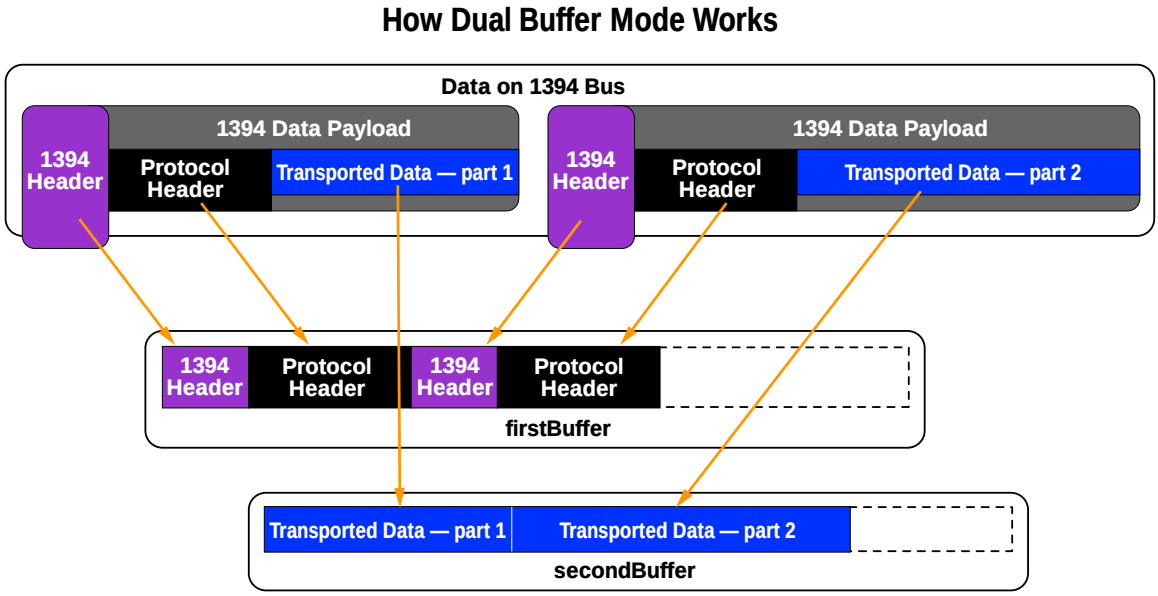

Figure 2

Since the buffer receiving the headers, firstBuffer, must have a length that is an exact multiple of the fixed size of the headers being placed in it, there is never a problem with the data from a 1394 packet exceeding the remaining space in the buffer. However, since the transported data is of variable length, it is possible for this data to exceed the space remaining in secondBuffer. In this case the DMA moves to the next descriptor and places the remaining data in the secondBuffer specified by this new descriptor in much the same fashion as Buffer Fill Mode. Figure 3 shows how this works.

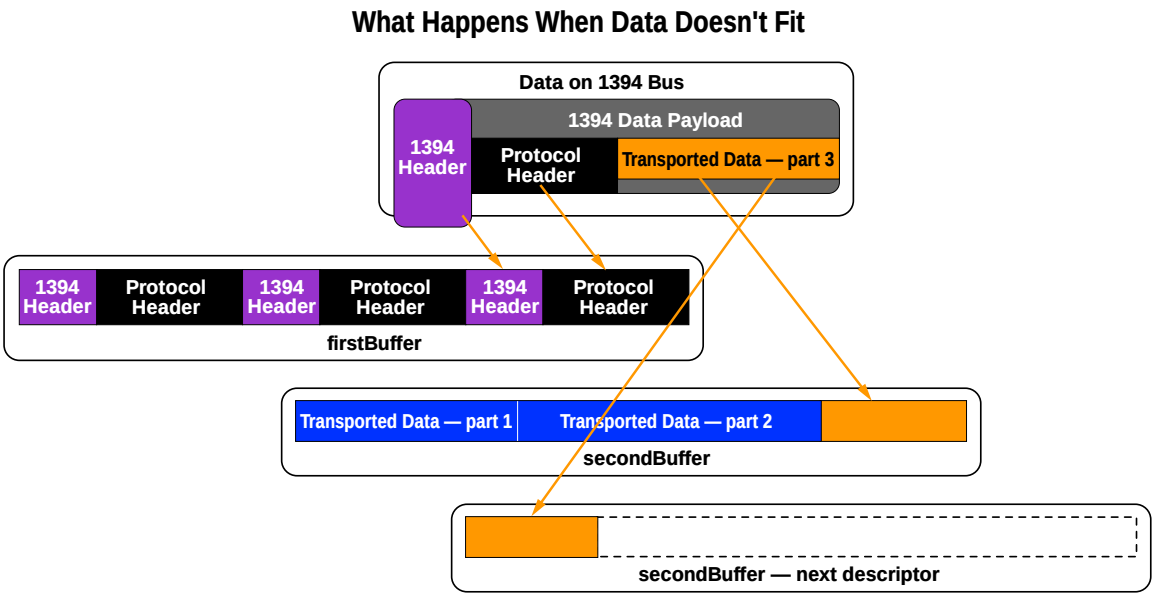

Figure 3

The DMA completes the descriptor either when firstBuffer is filled (exactly) or when secondBuffer is filled (exactly or when data overflows into next descriptor's secondBuffer). That is, the DMA completes the descriptor when either firstResCount or secondResCount is written as zero in the original descriptor.

The following restrictions apply:

- firstSize must be a multiple of 4 (that is, a number of quadlets).
- firstSize must be >=8 if 1394 headers are included (that is, isochHeader bit set).
- firstReqCount must be a multiple of firstSize (that is, holds an integral number of headers).
- firstBuffer must be on a quadlet boundary.
- Dual Buffer Mode cannot be used when multiChanMode set for the context.

If a 1394 packet is less than firstSize in length (including 1394 header and time stamp), then the short information will be written to firstBuffer. The next packet address will still be firstSize bytes from the start of the short packet header. Figure 4 shows this arrangement.

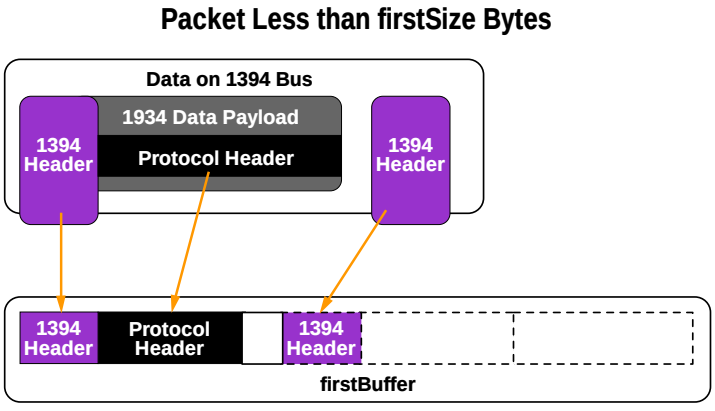

Figure 4

ack_data_error (Sections 5.4, 12.4, & 8.4.2.2)
There are problems with the way OHCI 1.0 uses and deals with ack_data_error:

- Asynchronous Transmit Response and Physical Response: They do retries of responses that received ack_data_error, but this violates IEEE 1394–1995 and IEEE 1394a and causes some interoperability problems.
- Asynchronous Receive: Responses for which ack_data_error was given are backed out, so software is not aware of a problem until a split transaction time-out occurs.

OHCI 1.1 fixes these problems as follows:

- Asynchronous Transmit: If ack_data_error is received, then complete descriptor with ack_data_error status (no retries).
- Physical Response: If ack_data_error is received, then pretend it was ack_complete.
- Asynchronous Receive: Don't send ack_data_error; use ack_busy_* instead (see IEEE 1394a, clause 10.9).

New Effects of busReset Event (Section 6.1)
In OHCI 1.0, software Compare/Swap access (see OHCI Section 5.5.1) is unsafe because a bus reset is possible just before this interface is used. This could cause software to store the wrong value in an IRM register. The Asynchronous and Physical filter registers (see OHCI Section 5.14) are also unsafe and could cause software to give access to an unsafe node or to block access from a safe node.

In OHCI 1.1, writes to the CSRControl, AsynchronousRequestFilterLo, AsynchronousRequestFilterHi, PhysicalRequestFilterLo, and PhysicalRequestFilterHi registers are ignored when the busReset event bit is set.

Power Management and ack_tardy
Three aspects of power management are new to 1394 Open HCI:

- PCI 1.1 Power Management Support
- LinkOn Port_Event notification from the PHY
- ack_tardy response

LinkOn Port_Event notification will be described as it applies to the various "D" states.

It is important to understand the definition for the word "should" as it is used in the 1394 Open HCI specification. "Should" implies a flexibility of choice with a *strong* preference for implementation. Note that the specification states PCI 1.1 level of power management *should* be incorporated into all 1394 Open HCI devices.

When PCI power management is implemented, four device states are required:

- D0_Uninitialized
- D0_Active
- $D3_{HOT}$
- $D3_{COLD}$

States D1 and D2 *should* be implemented.

D0_Uninitialized
When in the D0_Uninitialized state, the link is not asserting LPS to the PHY. A LinkOn signal from the PHY may occur and may persist until LPS is asserted. The LinkOn signal may be the result of a LinkOn PHY packet or the result of a Port_Event.

D0_Active
When in the D0_Active state, LPS is asserted from the link to the PHY. In the D0_Active state, PORT_EVENT notifications via PHY Status interrupt mechanisms are valid.

D1
When in the D1 state, the link continues to assert LPS to the PHY. An unmasked interrupt event will set PMCSR.PME_STS (a PME# may be generated if PME_EN ==1). An ack_tardy response (see Table A-11 in the specifications) is given for any unit access. The PCI, 1394 configuration, and GUID register contents are preserved. The OHCI will not attempt any host bus access while in D1.

D1 is the only state in which ack_tardy may be asserted.

When D1 is implemented, ack_tardy must be supported.

In implementations where it is not appropriate to include D1 state support, ack_tardy_enable bit shall be implemented as a permanent zero value.

D2
While in the D2 state, the link no longer asserts LPS to the PHY. All functional interrupt events are masked. A LinkOn PHY event will set PMCSR.PME_STS. (Note: a PME# will be generated if PM_EN ==1.) PCI configuration, GUID register, and PME context is preserved in D2. All 1394 configuration is lost.

$D3_{HOT}$
When in the $D3_{HOT}$ state, the link does not assert LPS to the PHY. All functional interrupt events are masked.

A LinkOn PHY event will set PMCSR.PME_STS. (Note: a PME# will be generated if PM_EN ==1.)

GUID register and PME context is preserved. PCI configuration and all 1394 configuration are lost.

$D3_{COLD}$
In the $D3_{COLD}$ state, the link is not asserting LPS to the PHY, and all device configuration and context is lost.

Out-of-Order Pipelining
In 1394 Open HCI 1.0, packet transmission is sequentially consistent (one packet follows another). Status is written back for the current packet before an attempt to transmit the next packet is made. Host bus error reporting is precise, and prefetching of packets is discouraged. So, what is the problem?

An inability to perform packet pipelining limits performance. Communication with a low-performance device will hinder interaction with a high-performance device, for example.

The 1394 Open HCI Revision 1.1 allows packet status to be written back in an out-of-order sequence. This has the effect of making host bus error reporting less precise (stating which packet is associated with the error, for example) and requiring special handling for event reporting.

Out-of-order pipelining is recommended for 1394 Open HCI 1.1 implementations.

Software Implications
To ensure sequential consistency, only a single packet may be enqueued at any one time.

Software is required to handle a dead context differently. The command pointer refers to the furthest processed descriptor block, and there is a need to scan context program and infer an ack_missing where the block was not updated.

Hardware Implications
Out-of-order pipelining requires hardware to implement multiple Asynchronous Transmit FIFOs and retry in the FIFO. Circulation of pointer/status implementation must be carefully done with retries from host memory.

Hardware requires careful design. Consult the 1394 Open HCI 1.1 specification for detailed descriptions of out-of-order pipelining.

Important Classification

RegAccessFail
The RegAccessFail register is used to detect when, for whatever reason, SCLK may have been intermittent and, more importantly, when PHY configuration registers may not contain values expected by the link layer. SCLK is controlled by the PHY and is enabled when the link asserts LPS. Initially, LPS is not active; therefore, SCLK is likely to be not active on initial start-up.

Since some registers in the OHCI device may be implemented in the SCLK domain, host accesses to OHCI registers may not function correctly if SCLK is missing.

In OHCI 1.0 there is no way to report a host bus access failure.

There is also an issue when the PHY power cycles separately from the link. This separate cycling may occur when a power source different from the link power source powers the PHY, as when the PHY is cable-powered and the link is system-powered, for example. (See Figure 5.) Not only will SCLK be lost, but initial PHY register configuration will be lost as well. In OHCI 1.0, there is no way to notify the link that the PHY has power cycled.

PHY Powered Separately from OHCI

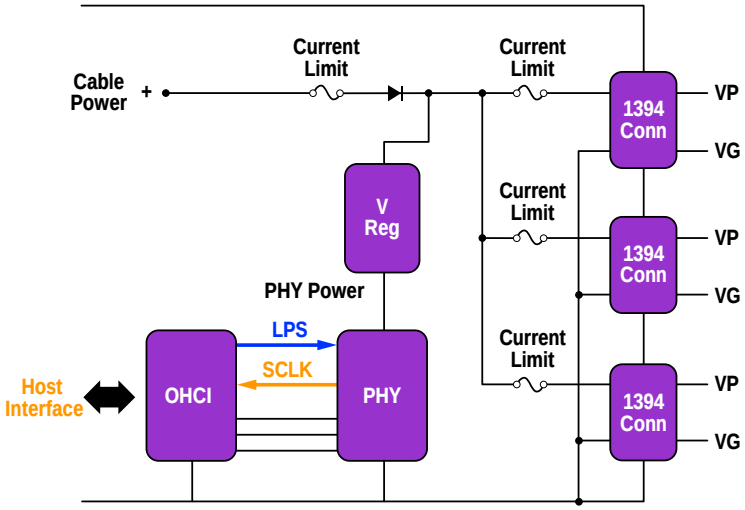

Figure 5

Other bus power problems can cause systems to lose SCLK or PHY register configuration. A temporary power short can blow a fuse, or the power provider can become detached from its power source. Figure 6 illustrates loss of SCLK.

Issues with Loss of SCLK Due to Bus Power Problems

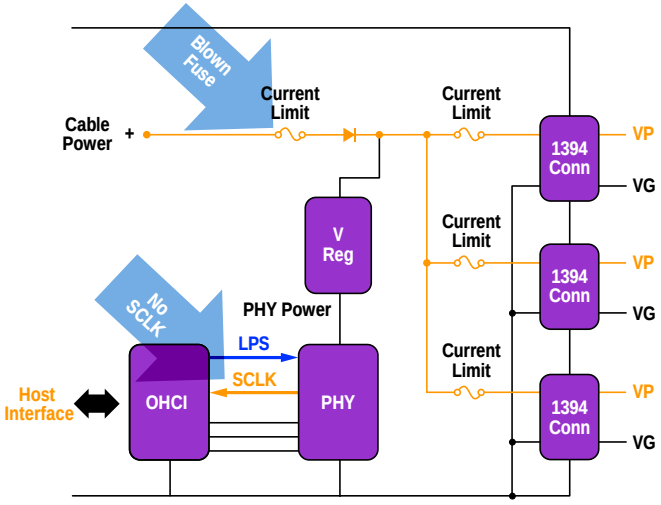

Figure 6

The way to solve the loss of SCLK issue is by providing notification for register access fail due to missing SCLK. The 1394 Open HCI Revision 1.1 specification creates a new interrupt event, ***IntEvent.regAccessFail, bit 18 rscu***. The characteristics of this event make certain there is no host bus error. A *regAccessFail* notification during a read yields an undefined value. A write has an undefined effect.

The occurrence of *regAccessFail* indicates that a 1394 Open HCI register access failed due to a missing SCLK clock signal from the PHY.

In paragraph 4.0 ("Register Addressing"), 1394 Open HCI Revision 1.1 provides a list of registers that are permitted to be implemented in the SCLK domain. All other registers must operate correctly regardless of the presence of SCLK.

Software Issues
Other than the *regAccessFailed* interrupt, there are no side effects.

Software must check for the *regAccessFail* interrupt after any host access(es) to SCLK domain registers.

- Any write must not be considered reliable.
- Any read data may not be considered valid.

Note: Software may wait until after the conclusion of a block of contiguous accesses.

AT PHY Packet Transmit
The 1394 Open HCI 1.0 specification does not restrict the transmission of AT PHY packets to two quadlets. This arrangement violates IEEE 1394a–2000 security requirements.

The solution provided in 1394 Open HCI 1.1 is to articulate that AT PHY packet transmission shall be restricted to two quadlets.

Changes to ITDMA
The 1394 Open HCI 1.1 specification makes four basic changes to ITDMA: Handling of FIFO Underrun, Isochronous Transmit Interrupt When Skipping, Handling Skip Processing Overflow, and Command Pointer Is Now Visible to Software.

FIFO Underrun
In 1394 Open HCI 1.0, the text describing how FIFO underrun works is quite ambiguous and does not state whether all of the next cycle is to be skipped or whether the remainder of the current cycle is to be skipped. The new behavior described in 1394 Open HCI 1.1 minimizes skipping and can be configured to, effectively, retry on FIFO underrun.

As Figure 7 shows, the lost counter is not incremented. Skip processing is to be performed immediately on the current descriptor block and the remaining blocks for that cycle. Normal processing resumes in the next cycle.

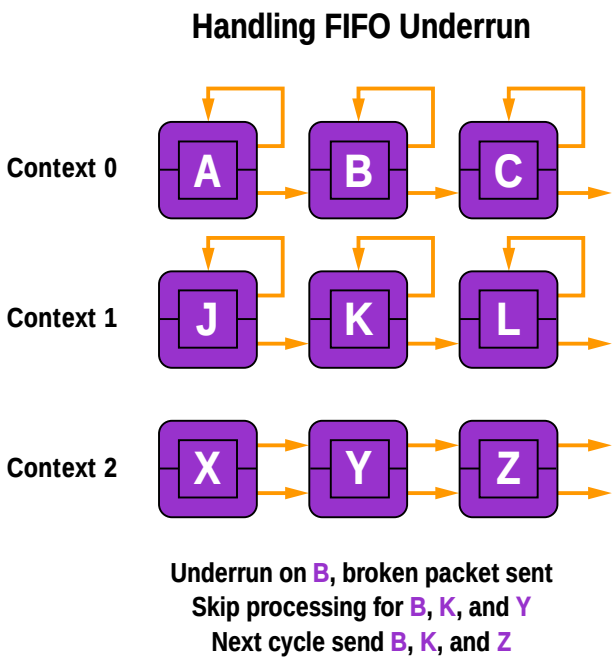

Figure 7

Isoch Tx During Skip Processing
A more reliable source of ITDMA interrupt is provided to software in 1394 Open HCI 1.1 because the first descriptor in a block may be configured to generate an Isoch Tx event when invoking skip processing.

Skip Processing Overflow
The 1394 Open HCI 1.0 specification failed to address a lost counter overflow condition. Without a way to report the overflow, ITDMA may not be able to perform the required amount of skip processing. In addition, implementations differ in lost counter size. Software is not able to detect cycle slips.

In 1394 Open HCI Revision 1.1, all implementations must handle at least three cycle skips (2-bit lost counter). Upon the occurrence of an overflow for each running ITDMA context:

- Set ContextControl.*dead*.
- Set IntEvent.unrecoverableError.
- ContextControl.*eventcode* == evt_timeout.

Command Pointer Visibility
In 1394 Open HCI 1.1, the CommandPtr is now valid for software to read when ITDMA context is active. The CommandPtr points to/into the descriptor block currently being processed by ITDMA context. This feature will help software to trace skip/branch paths.

Changes to Autonomous CSR Resources

CSR Resources and related changes in 1394 Open HCI C1.1, are described in:

- Section 5.5

—Bus Info Block
—Changed: Bus Manager and IRM
—Changed: Config ROM (first 1K)

- Section 5.8

—New: Initialization Registers

The following summarizes these changes:

- Atomic Config ROM update was not possible with OHCI 1.0. To support this, some registers work differently now.
- Automatic allocation (e.g., IP/1394 stream channel) is supported through three new registers.

The Config ROM experiences heavy access after bus reset. Other nodes scan GUID, Units, and so on. This is automated by OHCI:

- 5 quadlets from registers
- Remainder of 1K from memory

But there's a problem. Config "ROM" is not true ROM:

- Services (Units) come and go (application-specific protocols, proxy services for new devices)
- Software changes three registers: Header, Options, and Map Address

Atomic update is not possible, so inconsistent behavior could result. Also, by convention, Config ROM changes only on a bus reset, but software cannot synchronize a new Config ROM with a bus reset. The solution in OHCI 1.1:

- Software prepares a new Config ROM and indicates to hardware when it is ready.
- Hardware keeps using the old ROM until a bus reset activates new ROM.

This solves both problems: update is atomic, and update is synchronous with a bus reset.

How Config ROM Update Works

The following summarizes how this feature works:

- Config ROM mapping address (Section 5.5.6)

—Write has no immediate effect.
—Hardware continues to use the old value.
—Read returns the old value.

- Bus reset triggers update

—Hardware loads header (Section 5.5.2); options (Section 5.5.4).
—Hardware starts using new map address.
—New address becomes visible in register (Section 5.5.6).

- Map address (Section 5.5.6) has shadow

—Address is not directly accessible by software.
—Writing address (Section 5.5.6) arms update, but only if LinkEnable and BIBImageValid.
—Bus reset triggers update.

Config ROM Headers and Bus Options

The Config ROM Header is defined in Section 5.5.2. Bus options are defined in Section 5.5.4. Both are still read/write. However, if LinkEnable is set, software will not write. Both auto-update from memory on bus reset, as described in Section 5.5.6, but only if software wrote the map address (Section 5.5.6).

Software Strategy

Software uses two ROM buffers; the Config ROM Map (Section 5.5.6) points to one. Software writes new ROM, including new header and options, in an idle buffer. Software then sets a new Config ROM Map and causes a bus reset (Section 5.11). Hardware performs an atomic update following the bus reset.

Software can use two buffers, but hardware is not required to compare the buffers. Any write per Section 5.5.6 arms the update.

Hardware must ignore old requests. CSR requests in FIFO before a bus reset must not be answered after reset.

IRM Default Values
Invisible registers are defined in Section 5.5.1. On bus reset, IRM are reset to defaults:

- Bandwidth: 13'h1333
- Channels Hi/Lo: 32'hFFFF_FFFF

Owners must reallocate resources, but some resources may be reserved, for example, channel 31 for IEEE 1394a–2000.

New Registers
New registers are described in Section 5.8.

- Bus Management CSR Initialization

—Initial Bandwidth Available (0B0)
—Initial Channels Available Hi (0B4)
—Initial Channels Available Lo (0B8)

- Same defaults as before

—Bandwidth: 13'h1333
—Channels: 32'hFFFF_FFFF

- Adjustable defaults

—On bus reset, copy (Section 5.8) to (Section 5.1)
—(Section 5.8) is not changed by bus reset

This register scheme is efficient because fixed allocations are always automatic, and it's safe because a legacy device can't allocate reserved resource.

Physical Request Filter Registers
PhysicalRequestFilterHi/Lo (Section 5.14.2)

- physReqResource0 ... physReqResource62
- physReqResourceAllBuses … all non-local requests

The following was changed in 1394 Open HCI 1.1:

- physReqResourceAllBuses is no longer cleared by a bus reset
- All other bits are still cleared by a bus reset

CSR Compare/Swap—Compare/Swap access (Section 5.5.1) is unsafe as defined in OHCI 1.0. A bus reset is possible just before use, and software could store the wrong value.

Asynchronous Physical filters is also unsafe. Software could give access to unsafe node or block access from safe node.

The solution provided in 1394 OHCI 1.1 adds generation check; however, CSR Control register (Section 5.5.1) is not good enough. Therefore, the HCControl.busResetLockout bit has been added.

A bus reset sets the HCControl.busResetLockout bit. When set, host bus writes to CSR Compare/Swap (Section 5.5.1), Asynchronous request filter (Section 5.14.1), and Physical filter (Section 5.14.2) are ignored.

The HCControl.busResetLockout bit must be cleared by software after a bus reset.

SelfID Changes

Problems in 1394 Open HCI 1.0 include the following:

Non-persistent selfID Complete event—The selfID Complete event disappears with a subsequent bus reset, which can cause a spurious interrupt condition. The nselfIDComplete2 event stays set until cleared by software.

Race condition when updating self ID buffer—The text in 1394 Open HCI 1.0 is inadequate, implying a Self ID buffer update sequence that contains an inherent race condition, where quadlet 0 of selfID buffer is updated first. There is not enough information to guide implementers.

Width of selfIDSize field—Size in 1394 Open HCI 1.0 is too wide—4 K size, but 2 K aligned selfIDBufferPtr, which infers an adder.

The solution for SelfID context in 1394 Open HCI 1.1:

- Quadlet 0 of the selfID buffer is updated last (after all self ID packets).
- Generation number must be sampled and remembered for each self ID stream in a receive FIFO.
- SelfIDComplete event is raised after updating quadlet 0 of selfID buffer.

SelfIDSize field solution:

The size in 1394 Open HCI 1.1 is smaller, 2K. Building an address by concatenating offset with selfIDBufferPtr is acceptable.

More Info

The OHCI Promoters Group comprises Apple Computer, Inc.; Compaq Computer Corporation; Intel Corporation; Microsoft Corporation; National Semiconductor Corporation; Sun Microsystems, Inc.; and Texas Instruments, Inc.

To obtain a more detailed presentation of the changed and new content for 1394 Open HCI, access the complete 1394 Open Host Controller Interface Specification, Release 1.1. The specification is available from the Intel Developer Web site.

Support and directions for IEEE 1394 under Windows* operating systems is also available at Microsoft's hardware developer Web site.

Author Bios

Steve Bard is manager of the Mobile Interconnect Group within Intel Mobile Technologies Lab. Active in the development of 1394 technologies for the past three years, Steve was instrumental in driving energy conservation mechanisms—specifically, suspend/resume and standby/restore—into the 1394 Open HCI specification. Steve is an editor and major contributor to the 1394 TA Cable Power Distribution specification and Suspend/Resume Implementation Guidelines specification. He is also secretary to the IEEE P1394b Working Group.

John Fuller is program manager for External Buses, Microsoft Windows Core OS Device Drivers. He has been active in the development of 1394 technologies for the past 5-1/2 years, and instrumental in developing SBP-2 and RBC storage solutions. John is chair of the 1394 Trade Association's Energy Conservation Working Group. He is also task leader for the 1394 TA Architecture Working Group's Discovery and Enumeration Protocol specification.

Universal Plug and Play Connects the Home

Andrew Liu
Technical Marketing Manager
Intel Architecture Labs
Intel Corporation

Overview

Universal Plug and Play (UPnP) is an architecture for peer-to-peer network connectivity for PCs, appliances, and a variety of networked devices. It offers consumer value through compelling connected experiences, utilizing easy, affordable Internet access and connectivity—anywhere in the home. UPnP plays a key role in providing this connectivity by enabling command and control between devices on a network, independent of particular operating systems, programming languages, or media. In bringing simple, flexible, standards-based connectivity to consumer networks in the home and in small businesses, UPnP offers consumers easy device and service connectivity that doesn't require their intervention for network configuration, setup, or maintenance.

As a peer-to-peer connectivity model, UPnP supports zero-configuration networking and automatically discovers devices attempting to join a network. Utilizing Web-based IP, TCP, UDP, HTTP, and XML technology, a UPnP device can obtain the IP address for a device, announce its name, and convey capabilities upon request. Because a UPnP device may contain multiple devices and services, it can also detect the presence and capabilities of other devices or services on the network.

To ensure standardized authoring of specifications for UPnP devices, the Universal Plug and Play Forum was established by major corporations like Intel, which share a commitment to enabling connected devices in home and business environments. Formed in 1999, the Universal Plug and Play Forum now has over 140 member companies from consumer electronics, computing, home automation and security, home appliance, computer networking, and other related industries.

The five basic phases of UPnP are described in the sections of this article.

Device Discovery

In the discovery segment, depicted in Figure 1, clients find services, and services announce their presence to clients. UPnP utilizes Simple Service Discovery Protocol (SSDP), with a variant of HTTP that operates over multicast UDP for announcement and another variant of HTTP over unicast UDP for replies.

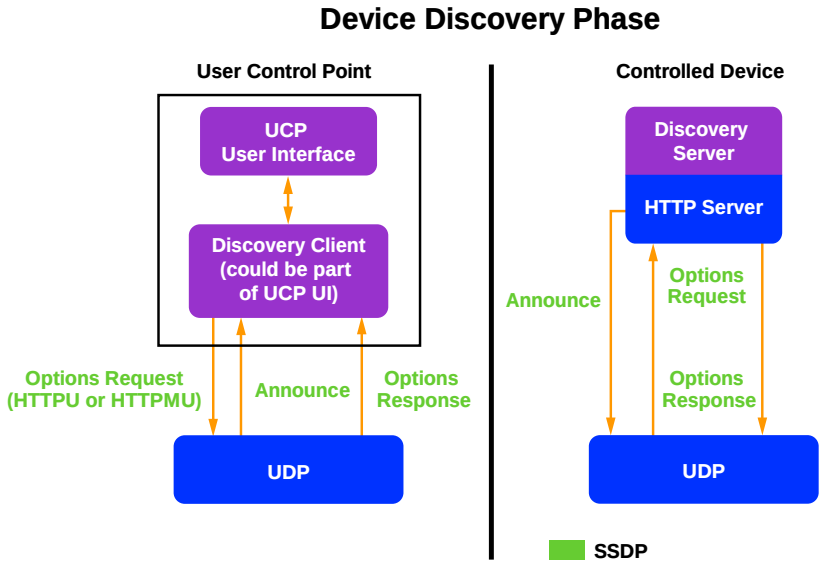

Figure 1

Clients multicast HTTP M-SEARCH discovery messages; devices that match the criteria specified by the client will respond with a URL that links to a device description document. If the client receives one or more favorable replies, the description phase is initiated.

If the client does not multicast any requests, servers can announce their device availability with a NOTIFY command. Clients receive the NOTIFY multicast and can then request the description document. Servers can also utilize the NOTIFY command to announce that their services are no longer available.

Device Description

Once the client has identified a service that may be of use, a description document is requested via HTTP over TCP, as illustrated in Figure 2. The client performs a standard HTTP GET command (similar to one that retrieves a Web page), and a complete description is sent to the client to enable a decision on whether or not the device is needed.

The description is an XML document that contains a range of information on the device, including manufacturer information and version, a list of services that the device supports, a list of embedded devices, and URLs to icons that can be used for the device. The description also provides information on how to control the device and how to request notification of status changes.

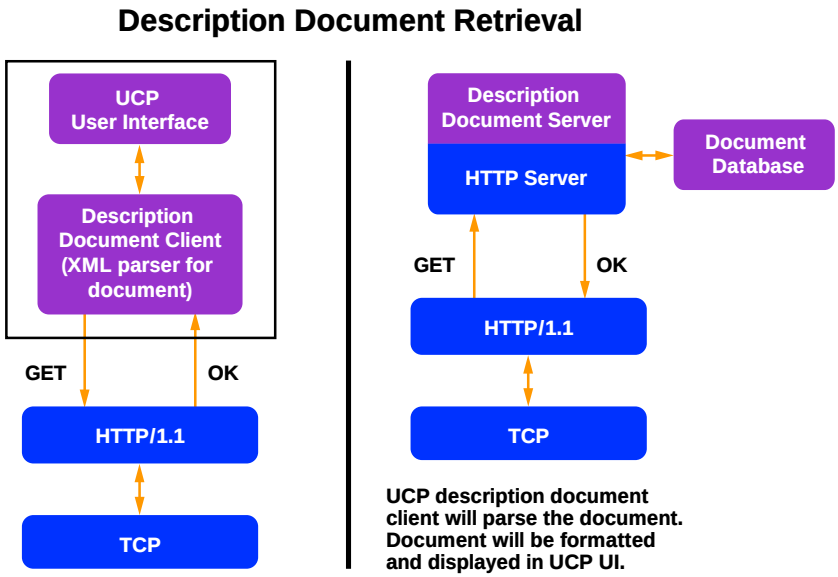

Figure 2

Presentation

Once the client has the description document, it may display the user interface provided by the device, use the services provided by the device, and receive events generated by the device. For devices that need or support a user interface for the device, the client can download an HTML document that provides the means of control or status.

The protocol for retrieving the usage (or presentation) document is the same as the one that retrieves the description document, HTTP over TCP. Not all devices have a presentation, and not all clients will
be able to display an advanced presentation document. However, if the presentation is available, the URL is contained in the description document and can be requested.

Events and Subscriptions

Once the client has discovered a device and retrieved the description and presentation, the client may want to be notified of status changes in the device. The URL for subscribing to event notification is contained in the description document, and event notifications will be sent to the client's event sink URL anytime the status of the service changes.

To subscribe or unsubscribe to event notification, an HTTP/TCP connection is made to the subscription URL from the description document, and a SUBSCRIBE or UNSUBSCRIBE request is issued. The client specifies an event sink URL where event notifications should be made, and events come via HTTP/TCP to the URL registered with the service. On

the server side, an event server waits for subscribe/unsubscribe requests using the General Event Notification Architecture (GENA) protocol, as shown in Figure 3.

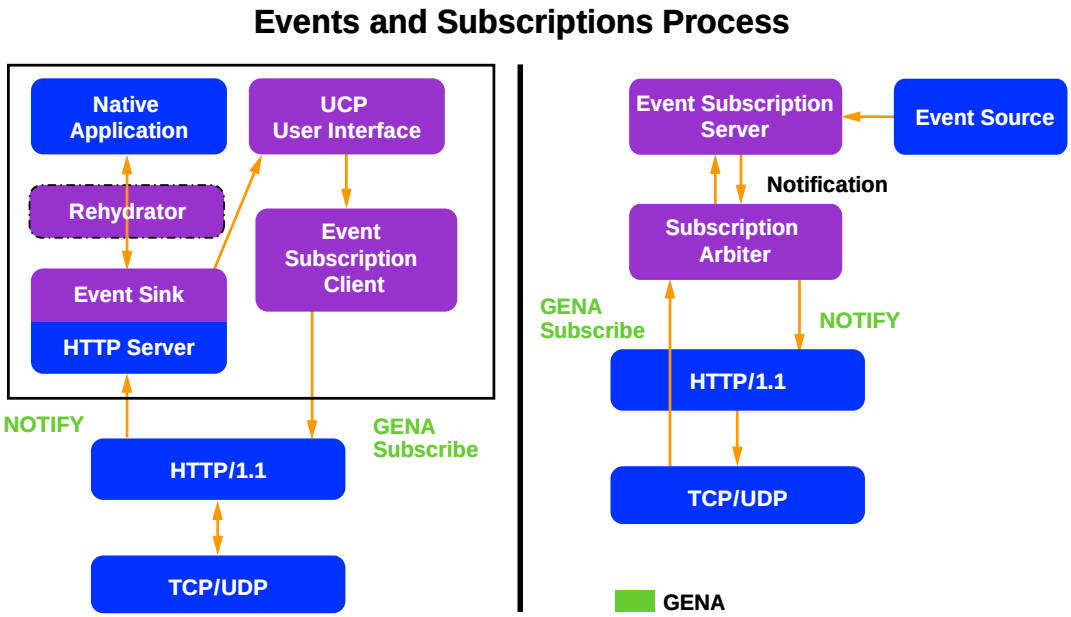

Figure 3

Control

After the client has discovered an interesting or useful device and retrieved the description document, it may wish to control one or more of the services contained in the device. The Simple Object Access Protocol (SOAP) allows a client to query or change elements in a service status table, using the POST or M-POST command sent over HTTP/TCP, as Figure 4 illustrates.

SOAP uses XML to specify what actions to take. The client creates the XML document from the commands found in the description XML and posts it to the control URL for the service specified in the description document. This allows the client to query for values or changes in the status table.

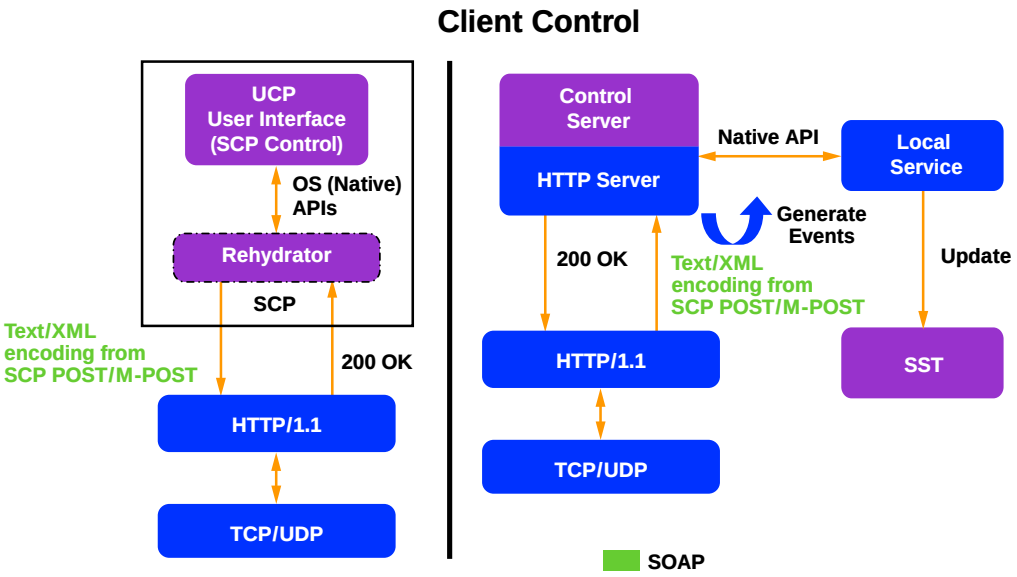

Figure 4

UPnP in Action

Let's say that a VCR (client) is to download an electronic program guide (EPG) from the Internet, using a dial-up modem (device) as the gateway.

First, the VCR must find the Internet Gateway by sending out a request. The Internet Gateway identifies itself by sending a URL to the description document. The VCR receives the URL, does an HTTP GET, and the document is returned. (See Figure 5.)

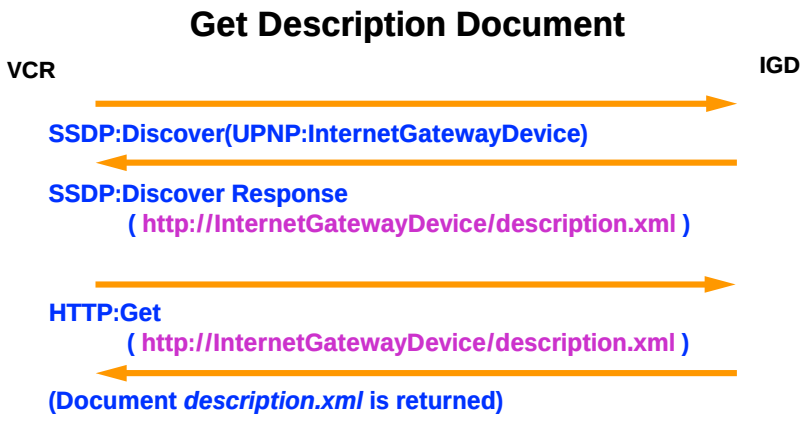

Figure 5

Next, the VCR parses the XML document; it discovers the variables, commands, and the location of the control and subscription URLs, and then requests the connection. (See Figure 6.) A configure and a request connection command are sent (HTTP POST SET), the connection is made, and a response is sent back to the VCR that the connection is made. The VCR uses the connection to retrieve the EPG. UPnP does not browse the Web site or initiate the download; UPnP is simply the protocol that enables device-to-device communication.

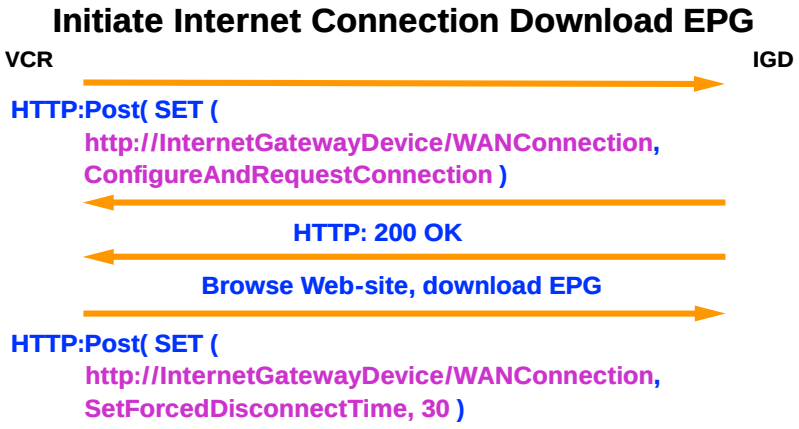

Figure 6

There are many ways to terminate the VCR's connection to the Internet. The VCR can request termination directly or the connection can be terminated automatically when the Internet Gateway detects inactivity. In another option, another device on the network (such as a clock) can initiate termination through an event notification when the allotted time has run out.

Summary

Universal Plug and Play (UPnP) is an easy-to-use, open standards-based, flexible platform. By taking advantage of Web technology in networked electronic devices, UPnP is making e-Home and e-Business a reality for the general population. This inexpensive, reliable network solution makes it possible for small computing and other devices to remotely control appliances; enable transfer of digital audio, video, and still images between devices; and share information among devices and the World Wide Web.

In five simple steps, utilizing clear command sequences, UPnP offers easy-to-use device manageability for networked devices. This benefits appliance users and information users alike. Universal Plug and Play takes existing standards, technologies, and knowledge, and applies them to deliver new opportunities and remove barriers in the networked world of the near
future.

More Info

For more information, please visit the Universal Plug and Play Forum Web site or visit the Connected.Home Initiative area of the Intel Architecture Labs Web site.

You can find the Universal Plug and Play Marketing Requirements Document, Version 0.5, March 20, 2000 in the resources area of the same site.

Author Bio

Andrew Liu is a technology marketing manager in the Intel Architecture Labs (IAL) and is responsible for Home Networking Infrastructure technologies, including residential gateways, Universal Plug and Play, and wireless technologies. He has been with Intel since 1993. Previously, Andrew worked with Intel's Anywhere in the Home, scalable platforms, and Internet media technologies. He has presented at the Intel Architecture Conference, ISPCON, the Universal Plug and Play Forum, and WinHEC. Andrew holds a B.S. degree in Electrical Engineering and Computer Science from the University of California at Berkeley, and a Master of Electrical Engineering from Cornell University.

Networking & Communications

Intel® 82562 PLC Enables Chipset LAN Solutions

Kevin Cline
Product Marketing Manager
Network Communications Group
Intel Corporation

Overview

Integration of communications on the desktop and between multiple desktops is, by definition, what networking is all about. By achieving increased levels of integration, end users realize greater cost savings while enhancing communication capabilities within the network. Until recently, having the highest integration of capabilities, as well as highest quality connections, meant paying a higher premium for networking solutions over less expensive, less integrated, and typically less reliable networking solutions.

Intel has increased the baseline for quality and management capabilities by providing the next significant innovation in networking technology—integrating LAN solutions into a platform chipset. This is achieved with the new, second-generation I/O Controller Hub (ICH2) found in Intel® 815E and 820E chipsets, the companion 82562 Platform LAN Connect (PLC) and associated SW. In doing so, Intel has provided a reliable solution for networking designs that targets the value market without compromising quality or performance.

Solutions for the Value Market

There have traditionally been two extremes in LAN communications: high-priced solutions with high quality, or low-cost solutions with reduced quality. Now, systems using the Platform LAN Connect (PLC) in conjunction with an ICH2-based chipset offer high-quality LAN connectivity at a low cost. The ICH2 solution also utilizes Intel® SingleDriver™ technology, which simplifies network complexity and increases the ease of deployment. This is achieved by allowing the network administrator to install Performance and Value connections in their network (or upgrade from Value to Performance as the network needs increase) with the same driver and installation utilities.

The ICH2 LAN solution is a 32-bit PCI solution that features enhanced scatter-gather bus mastering capabilities, allowing the LAN solution to perform high-speed data transfers over the PCI bus. These bus mastering capabilities enable the component to process high-level commands and perform multiple operations, minimizing CPU usage. Two large transmit and receive FIFOs are also included in the architecture to enhance overall performance.

The Intel® 82562ET and EM are highly integrated PLC devices, combining 10BASE-T and 100BASE-TX physical layer capabilities and power management capabilities. This provides a core ingredient to enable integrated network connectivity in Intel® ICH2-based platforms. Based on Intel's second generation of fully integrated 10BASE-T/100BASE-TX solutions, the integrated LAN solution has the stability and reliability engineers look for in designing applications.

While the EM version enables additional enterprise management capabilities, both the ET and EM solutions support a single interface that is fully compliant with the IEEE 802.3/802.3μ standard. Additionally, both support advanced power management capabilities like reduced power consumption when the network link is lost or disabled.

The ICH2 enables three LAN options, with a single LAN driver for all three solutions. The ICH2, software, and 82562ET/EM PLC support basic and managed 10/100-Mbps Ethernet and are compatible with Intel® PRO/100 adapters. A different 82562EH PLC, designed to Home PNA 1.0 specifications, enables phone wire Ethernet and compatibility with Intel® AnyPoint™ networking adapters.

The Media Access functionality in the ICH2 is connected to the PLC via a dedicated data path, the LAN Connect Interface. This approach offers several advantages over existing interfaces:

- A lower pin count than standard interfaces offers increased signal integrity, easier routing, and smaller component packaging for greater value to the system designer.

- By lowering the signal speed from the average 125 MHz offered by standard interfaces of 50 to 60 MHz while maintaining 100 Mbps throughput, motherboard space can be better utilized and noise is reduced, allowing for easier routing and longer traces.

Power savings through advanced management capabilities built into the connection allows networks to be cost-optimized in real time based on the overall network traffic. By adding the Low Power mode and Link sensing/status capabilities, a network can be power-optimized to the specific activity level of the users within that network.

A Variety of Form Factors

For added flexibility, target applications include (but are not limited to) LAN on Motherboard (LOM) and a Communication and Networking Riser (CNR) card for the Value segment.

LOM offers easier layout options, quick design, guaranteed interoperability, and the lowest cost option for both Value and Performance platforms.

The CNR card allows multiple configurations on a single card, and is smaller than a regular PCI adapter card. The riser requires a special CNR board connector that is about half the length of a shortened PCI slot. In the Intel integrated LAN solution, the LAN connect interface runs to the CNR card and the Platform LAN connect chip is mounted on the CNR card. Along with a LAN interface, the CNR can offer USB and SMBus support, as well as an AC ‘97 interface for audio and modem.

The introduction of the ICH2 also provides an effective solution for mobile users in the MiniPCI form factor. Designed to accommodate the compact dimensions of a laptop by mounting parallel to the motherboard (instead of vertically), the MiniPCI card can route the LAN Connect Interface, giving mobile users the same low-cost flexibility that the CNR offers.

Improving on a Connection

Intel’s integrated LAN solution not only provides basic LAN connectivity, but also delivers cost and performance benefits unique in the industry. Intel’s goal in creating this LAN solution was to improve the quality and capabilities provided as a baseline in this segment of the market.

Two primary objectives were identified: improving the software baseline and minimizing hardware design constraints and power usage. For the first objective, the focus was to use the high-quality and high-performance software capabilities that already existed in Intel® LAN solutions while maintaining reliability and responsiveness in software offerings. In the second objective, Intel sought ways to increase power management and minimize the number of signals, noise, and routing restrictions.

This resulted in the following benefits:
1) Intel® SingleDriver technology is available across all Intel® desktop solutions, from Performance to Value. This allows IT managers to use one high-quality driver for a variety of user capabilities and solutions. SingleDriver technology lets an IT manager tailor solutions to user needs without sacrificing simplicity in installation and maintenance. This optimizes IT budgets by only providing what is necessary for each individual user. Unlike competitive solutions, SingleDriver technology supports the ability to add new users with different needs or upgrade existing users while still using the same driver.

2) The external physical layer component and LAN Connect Interface provide a significant reduction in signals as well as signal frequency over standard interfaces, thus achieving improvements to the component size and overall use by the end system designer. Additionally, power management improvements supplement the basic capabilities in standard interfaces.

3) By separating the analog and power management functionality from the core chipset, Intel provides a wide variety of options as cost effectively as possible to the chipset consumer (no LAN, HPNA 1.0, 10/100 LAN, Managed 10/100 LAN).

Summary

If progress in network connectivity can be defined as "achieving higher degrees of integration," then the introduction of Intel's second-generation ICH2 represents a significant step forward. The Intel integrated LAN solution, including the ICH2, PLC, and associated software, delivers by providing a reliable, high-performance network connection at a significantly reduced cost.

Providing a Value-segment solution for networking designs is critical to meeting the growing segmentation needs of the computer industry. With the integrated chipset and the Platform LAN Connect, Intel has driven the next technology transition in networking.

More Info

You can find more in-depth information on Intel's Integrated LAN Solution at the Developer Web site:

- Technical information on the 82562ET/EM Platform LAN Connect
- Technical information on the 82562EH Platform LAN Connect
- Communication and Networking Riser Card
- Technical information on the 815E chipset
- Technical information on the 820E chipset

For a related article, focusing on the ICH2 features in the 820E and 815E chipsets, read the Desktop article, "New Chipsets Deliver Performance and Flexibility," in this issue of *Intel Developer Update* magazine.

Author Bio

Kevin Cline is a product manager in the Network Communication Group and is responsible for networking solutions integrated into Intel chipsets. He joined Intel in 1994 as an applications engineering manager and has been active in the definition and support of four generations of Intel 10/100 Ethernet silicon. Kevin won an Intel Achievement Award in 1999 for his work on manageability in networking products, and he has two patents to his credit. He currently serves as Chairman of the Pre-OS Working Group of the Distributed Management Task Force. A graduate of California State University, Sacramento, Kevin has a B.S. in Computer Engineering.

Servers

Intel® 80303 I/O Processor Boosts PCI Bandwidth

Lance K. Packer
Senior Product Marketing Engineer
I/O Products Division
Intel Corporation

Overview

As new Internet and storage applications continue to make growing demands on servers, improving I/O throughput and balanced system performance have assumed greater importance for developers. The Intel® I/O processor family is designed to help meet these requirements. The Intel® 80303 I/O processor is Intel's third-generation I/O processor. It continues Intel's vision of improving I/O performance on current and future server platforms, including servers based on the Intel® Pentium® III Xeon™, and Itanium™ processors.

The "80303" designation inaugurates a new naming convention for Intel's I/O processor roadmap; it is the first I/O processor name that does not emphasize the processor core. While the 80303 I/O processor is based on the 100-MHz 80960 JT core, with large instruction and data caches, its performance improvements are due to the on-chip integration of architectural features designed to enhance data flow, including a 64-bit, 66-MHz PCI-PCI bridge for improved access to PCI-enabled peripheral products.

The 80303 I/O processor meets five principal design goals:

- Maximize bus bandwidth by increasing PCI and internal bus speeds and by utilizing deeper queues.
- Isolate data traffic on the PCI buses.
- Increase SDRAM memory speed.
- Minimize interrupts to host CPU by handling all I/O processing.
- Improve hardware support for RAID 3 and 5.

Pumping up Bandwidth

Based on input from customers and experience with previous-generation Intel® i960® I/O processor designs, Intel has created its highest performance integrated I/O processor to date.

The 80303 is Intel's first I/O processor with an integrated 66-MHz, 64-bit PCI-to-PCI bridge. PCI bandwidth has been doubled compared with the previous-generation I/O processor.

The 80303 internal bus speed is enhanced to run at 100 MHz, supporting internal bandwidth up to 800 Mbytes/sec.

The integrated memory controller now supports up to 100-MHz SDRAM with up to 512 Mbytes of memory, which is four times the memory supported by the previous generation.

Six secondary PCI output clocks, 4 SDRAM output clocks, and 8 general purpose I/O pins were added to help reduce board cost and simplify designs. In addition, the 80303 features an improved ball map to facilitate board layout.

Like its immediate predecessors—the i960 RM/RN I/O processors—the 80303 I/O processor features an integrated programmable Application Accelerator Unit that enables hardware-based XOR calculations for RAID levels 3 and 5, two Address Translation Units, and three DMA (Direct Memory Access) controllers between the PCI buses and local memory.

Applications

The 80303 I/O processor provides a highly integrated, cost-effective way to add intelligence to any server subsystems that require fast movement of large amounts of I/O data. Examples include server-attached storage (RAID, SCSI, Fibre Channel, and Storage Area Networks), Networking (LAN, ATM, Ethernet, switches, and routers), workstations, and various embedded applications. Because of its faster bus speeds, the 80303 I/O processor can support performance gains in RAID applications (RAID 0) of up to 25 percent, compared with previous-generation I/O processors.

Using the 80303 I/O processor as a dedicated coprocessor provides the additional benefit of balancing system-level performance in servers. This is because the 80303 I/O processor reduces host CPU utilization by servicing the I/O requests directly. The result can be an overall increase in server application performance. Off-loading I/O tasks from the host CPU to the I/O processor also helps servers scale to meet additional storage and networking requirements.

Bus Speed Improvements

The integrated 66-MHz PCI-to-PCI bridge supports PCI bus bandwidth up to 528 Mbytes/sec. This bridge maintains backward-compatibility with 33-MHz PCI. The PCI buses can both be run at 33 MHz for loading reasons or if necessary, the primary PCI bus can run at 66 MHz and the secondary bus at 33 MHz to support legacy 33-MHz components. This provides the system designer with the flexibility to determine the best configuration for a specific application. In addition, the 80303 I/O processor PCI pins can be either 3.3V or 5V, commonly called “Universal PCI.” This allows the user the option of using the processor in either voltage environment.

The internal bus bandwidth of 800 Mbytes/sec. represents a 33 percent increase over the previous generation of I/O processors. Not only do all blocks attached to the internal bus benefit from the added bandwidth, but the Memory Controller Unit can now support 100-MHz SDRAM at full speed.

Software Tools and Support

Intel's IQ80303 QUICKval kit, based on the 80303 I/O processor, provides a platform for the rapid and cost-effective development of I/O software, including basic device drivers. The kit supports network operating system-to-driver independence and enables developers to implement multiple I/O software applications.

The kit includes a comprehensive set of software development tools, including Intel® CTOOLS and Wind River Tornado* tools, a real-time operating system (Wind River IxWorks*), and a JTAG emulator. It also includes related documentation that can help to further accelerate development time.

The CTOOLS software development suite contains a profile-driven C/C++ compiler, assembler, linker, runtime libraries, debugger, monitor, software development utilities, and online documentation. The compiler reduces the code-tuning effort and increases application performance by providing code execution profiling and instruction scheduling optimizations for the i960 processor architecture. IxWorks, coupled with Tornado for I_2O*, provides a powerful, general-purpose real-time operating system with clearly defined application program interfaces (APIs).

One other important note: compatibility with existing i960 processor-based code helps protect developer investments in existing software.

Summary

With the expansion of the Internet, networked computers are processing more data than ever before, and server I/O has become a pivotal performance issue for developers. Intel has created a family of I/O processors that help balance server performance while helping to relieve data bottlenecks.

Intel’s IOP303 I/O processor is an intelligent I/O subsystem on a chip, including an integrated 64-bit PCI-to-PCI bridge. Its tight integration reduces the number of components necessary for the system design; its system cost and board real estate requirements are lower than those of multi-chip solutions.

The 80303 I/O processor further optimizes I/O performance and data flow with an integrated 100-MHz internal bus and support for 100-MHz SDRAM, in addition to other tightly integrated features. To further speed development, the 80303 I/O processor is supported by more than 200 development tools for the i960 processor family, in addition to a comprehensive software development platform.

With the introduction of the 80303 I/O processor, Intel is providing a new cost-effective building block with the enhanced performance needed to support the growing I/O requirements of server-attached storage and networking applications.

More Info

Visit the design area of the Intel Developer Web site for the following information:

- 80303 I/O Processor Product and Evaluation Board Brief
- 80303 I/O Processor Datasheet
- 80303 I/O Processor Specification Update
- 80303 I/O Processor Developer's Manual
- 80303 I/O Processor Design Guide
- IQ80303 Evaluation Board Manual

Author Bio

Lance Packer is a senior product marketing engineer in Intel's I/O Products Division. He is responsible for development, launch, and support of Intel's family of I/O processors. Lance joined Intel in 1998, bringing over 12 years of industry experience with him. He holds a B.S.E.E. from Utah State University.

Software

Intel® Developer Services Enables Better Solutions

Melissa Laird
Director, Operations
Solutions Enabling Group
Intel Corporation

Overview

Intel recently launched a new Web site aimed at meeting the specific needs of software and Internet developers looking to deliver high-performance, cost-effective software solutions for Intel® Architecture (IA) platforms. The Intel® Developer Services Web site provides developers with resources, tools, and training from Intel and other industry experts that enable them to take advantage of the latest Intel technology.

Intel Developer Services

For years Intel has worked with individual tools, applications, and operating systems developers to ensure their products ran well on the latest Intel architecture, but this effort has been limited to the number of accounts an individual in the field could adequately support. To augment that effort, Intel is delivering a comprehensive program to assist developers worldwide who want to get to market faster with the best applications and solutions.

Intel Developer Services was created to meet the needs of a broad range of software and Internet developers, and use the Web to deliver a rich set of resources to them. It not only provides the technical information and support developers need to deliver great new products faster, it will also provide marketing support to help those products succeed in the marketplace. Intel is inviting software and Internet developers to become members of the program and begin accessing the Web site now. Over the next several months additional features and member benefits will be added.

Structuring the Site

Through research conducted over the past six months, Intel learned from developers what they value most in a developer support site, from downloadable libraries to Web-based training on emerging technologies. We also learned what they would value most from Intel, including unbiased development information on the range of operating systems now available for IA.

The Web site is organized both by topic and by industry segment to serve the varying needs of visitors. For example, an e-Business solutions provider can navigate through the complete breadth of elements required for a solution from high-level applications to low-level drivers. An application developer meanwhile can go directly to the performance-tuning tool he or she is looking for.

Key areas of the site focus on tools, training, and services, three areas developers told us they care about the most.

A comprehensive Tools section includes libraries, toolkits, information on browser compatibility, testing, performance optimization, hosting, and solution integration for Internet and e-Business developers. For application developers, the site offers access to hardware development vehicles, software development kits, libraries and engines, and information on multiprocessing, multithreading, compiling, and debugging.

The Training section features live Web casts on a range of topics. Recent Web casts include “Effectively Developing E-Commerce Web Sites on IA,” “How to Support Customers Quickly and Efficiently via the Web,” “Industry Enabling for IA-64,” and “Getting Software Ready for 64-Bit (IA-64) Implementation.” More than a dozen self-paced classes are also available, with more being added each month.

Developer Services includes application forms to visit an Application Solution Center or e-Business Solution Center, where Intel technical support people work directly with developers to expedite their time-to-market with optimized and well-tuned products. These Centers are state-of-the-art labs equipped with the latest Intel® hardware, the most advanced Intel® performance tools, and a staff of highly skilled and trained software performance engineers. Intel Solution Centers offer a lab environment where developers can tune their IA-based applications and port to the new IA-64 architecture.

Throughout the site, numerous case studies provide real-world examples of successful solutions. And of course, developers will also be able to access the latest technical information about new Intel® microprocessors, embedded systems, and Flash solutions available throughout Intel's family of Web sites.

Another feature of the site is Intel® Developer Services News, an electronic newsletter published twice a month to notify members of new information and features available on the site. The newsletter is distributed via e-mail to program members free of charge and is available in HTML or text format.

Summary

The Intel Developer Services Web site provides the latest resources and information developers need to deliver faster, richer software solutions to the marketplace. In addition, the site offers developers the unique perspective of the company most familiar with the Intel Architecture and related innovations.

Intel has launched the Developer Services site as a major new commitment to developers. Its goal is to provide timely and useful information on software, Internet, and e-Business solutions to help developers create high-performance and cost-effective applications and solutions.

More Info

To become a member, register at the Intel Developer Services Web site.

Author Bio

Melissa Laird is director of operations for Intel's Solutions Enabling Group, where she is responsible for delivering a comprehensive set of services to facilitate the growth and success of IA software and Internet developers and solution providers worldwide. She joined Intel in 1983 and has held numerous positions in marketing and manufacturing, including co-general manager for the PC networking adapter business, marketing manager for the Intel Architecture Labs, manufacturing manager for the Enterprise Server Group, and strategic materials manager for Corporate Materials. Melissa holds a B.S. in Mathematics and Computer Science and an M.B.A.

—End of Intel Developer Update Magazine Issue 10—